Data-Driven Allocation of Preventive Care
With Application to Diabetes Mellitus Type II

Mathias Kraus

FAU Erlangen-Nürnberg, Lange Gasse 20, 90403 Nürnberg, Germany, mathias.kraus@fau.de,

Stefan Feuerriegel

Munich Center for Machine Learning & LMU Munich, Geschwister-Scholl-Platz 1, 80539 Munich, Germany, feuerriegel@lmu.de,

Maytal Saar-Tsechansky

Department of Information, Risk and Operations Management, McCombs School of Business, University of Texas at Austin, 2110 Speedway, B6000 Austin, TX 78712, USA, maytal@mail.utexas.edu,

Problem Definition. Increasing costs of healthcare highlight the importance of effective disease prevention. However, decision models for allocating preventive care are lacking.

Methodology/Results. In this paper, we develop a data-driven decision model for determining a cost-effective allocation of preventive treatments to patients at risk. Specifically, we combine counterfactual inference, machine learning, and optimization techniques to build a scalable decision model that can exploit high-dimensional medical data, such as the data found in modern electronic health records. Our decision model is evaluated based on electronic health records from 89,191 prediabetic patients. We compare the allocation of preventive treatments (*metformin*) prescribed by our data-driven decision model with that of current practice. We find that if our approach is applied to the U.S. population, it can yield annual savings of $1.1 billion. Finally, we analyze the cost-effectiveness under varying budget levels.

Managerial Implications. Our work supports decision-making in health management, with the goal of achieving effective disease prevention at lower costs. Importantly, our decision model is generic and can thus be used for effective allocation of preventive care for other preventable diseases.

Key words: preventive care; disease prevention; healthcare analytics; machine learning

History:

1. Introduction

Millions of people die each year of diseases that could have been prevented. Common examples of preventable diseases include certain infectious diseases, such as human immunodeficiency virus (HIV); respiratory diseases, such as pneumonia; and cardiovascular diseases, such as diabetes mellitus. Over one quarter of all deaths in OECD countries are premature and could have been avoided through better preventive efforts (OECD 2019). Furthermore, the prevalence of preventable diseases is likely to rise

as a result of current lifestyle choices (World Health Organization 2018c). Preventable diseases reduce individuals' quality of life and also entail substantial costs. The cost of treating preventable diseases in the United States now totals $730.4 billion annually, which corresponds to 27.0 % of total healthcare spending (Bolnick et al. 2020).

Preventive care aims at managing risk factors before the onset of a disease (World Health Organization 2018b). A prominent example of preventive care for diabetes mellitus type II is *metformin*, which is a preventive drug aimed at controlling impaired glucose tolerance (Knowler et al. 2002). However, many effective preventive treatments are often costly, and the effectiveness varies by risk factors (Zhou et al. 2020). Consequently, healthcare managers face the challenge of cost-effective allocation of potentially limited or costly preventive treatments to patients. Indeed, with the growing importance of preventive care, effective allocation is essential in improving the quality of lives of billions of people worldwide. However, decision models for informing allocations in preventive care are lacking.

Prior research has considered decision support for health management with the goal of achieving effective healthcare at lower costs. For instance, research has proposed optimizing periodic checkups (e. g., Ayvaci et al. 2012, 2017, Kamalzadeh et al. 2021, Liu et al. 2018), as well as the selection of treatment designs (e. g., Ibrahim et al. 2016). However, these works use generative models that are computationally prohibitive in the analysis of high-dimensional data to accurately predict future disease onsets. In contrast, discriminative machine learning models effectively scale to large datasets involving millions of observations and thousands of variables (e. g., Choi et al. 2017), which is common in modern healthcare settings. However, these models have not been adapted for allocation problems among patient populations in preventive care; this adaption is the objective of our work.

Our work contributes to health management by proposing a data-driven decision model that provides decision support for allocation problems in preventive care. For this aim, we combine counterfactual inference, machine learning, and optimization techniques to exploit high-dimensional medical data from electronic health records (EHRs). To our knowledge, this work contributes along three paths: (i) we propose a data-driven decision model for allocating preventive treatments; (ii) we demonstrate the value of scalable machine learning from rich, high-dimensional data in electronic health records for preventive care; and (iii) we establish the (cost-)effectiveness of data-driven allocations.

We evaluate our decision model for allocating preventive treatments for the case of diabetes mellitus type II. Diabetes is a chronic disease that can significantly reduce patients' quality of life and that comes with substantial long-term costs to treat its many complications (Lee et al. 2018). For our evaluation, we use EHRs from 89,191 patients with a prediabetic condition and analyze the allocation of preventive care in a dynamic setting spanning a period of 10 years. A simple, yet effective preventive

treatment for diabetes type II is *metformin*. However, applying *metformin* comes at a cost, thus raising the question of what a cost-effective allocation is in practice. We compare our decision model for preventive care allocation against a clinical baseline that stratifies preventive care according to the Framingham diabetes risk score. We find that our decision model prevents over 25 % more cases of diabetes mellitus type II. Further, we find that, if our approach is applied to the U.S. population, it can yield annual savings of $1.1 billion because of a more effective allocation.

Our decision model is directly relevant for healthcare organizations aimed at allocating preventive care under resource constraints. This is the case for vertically integrated healthcare organizations (IHOs) worldwide, where hospitals and health insurances exchange patient data with the aim of achieving better coordination, efficiency, and quality of care for patients (World Health Organization 2016). Examples of countries with IHOs are Israel, Singapore, many European countries, and partially the U.S. through health maintenance organizations (World Health Organization 2016).

Our work promises to have a direct effect on both practice and research. For healthcare practice, our research demonstrates that using a data-driven decision model based on scalable machine learning and EHRs can substantially improve the performance of preventive care programs and reduce the risk of disease onset. As such, our decision model supports health management in primary preventive care by offering a cost-effective path toward reduced health-related spending while improving patients' outcomes and quality of life. For research, our new decision model, incorporating counterfactual inference, offers a widely applicable solution for allocation problems in preventive care and for other allocation problems that require data-driven approaches in management science.

The remainder of our paper is structured as follows. Section 2 provides an overview of preventive care, revealing a scarcity of decision models for this purpose. Section 3 details our research setting, which focuses on allocating preventive treatments to patients at risk of developing diabetes mellitus type II. In Section 4, we develop our data-driven decision model, combining counterfactual inference, machine learning, and optimization techniques to determine a cost-effective allocation of preventive care to patients at risk. We report our empirical results in Section 5. Finally, we discuss implications of our work for health management in Section 6.

2.   Background

Our work builds on literature in the fields of preventive care, machine learning, counterfactual inference, and resource allocation in health management. We review this work in the following sections.

2.1.   Preventive Care

Preventive care refers to measures taken to reduce people's risk of developing a disease or suffering from a severe course of illness. Preventive care can be divided into three categories: primary prevention,

secondary prevention, and tertiary prevention (Goetzel 2009). Primary prevention aims at averting the onset of a disease altogether. Examples include vaccines, smoking cessation, weight loss, and other measures that address risk factors through the promotion of good health (World Health Organization 2018c). Secondary prevention focuses on the early stages of a disease after its onset. Here, the objective is to prevent a disease from developing into a critical or acute condition. Such approaches include screening (e.g., mammography) and other forms of early diagnosis. Tertiary prevention targets the progression associated with the disease. Here, the aim is to reduce the severity of symptoms. Examples include mental, physical, and social rehabilitation, such as in cardiac rehabilitation programs. In this work, we focus on primary prevention—specifically, on preventing patients from developing diabetes mellitus type II.

Preventive care can be a valuable economic investment; when it is provided early to patients, it can reduce the need for more expensive treatment after the onset of an illness (World Health Organization 2018b). Of note, the main aim in preventive care is to improve the health benefit for patients. Hence, optimization is measured against the overall health outcomes relative to costs, rather than against cost minimization alone, for ethical and legal reasons (see Gold et al. (1996), National Institute for Health and Care Excellence (2015), and World Health Organization (2018b)). However, not all patients provided with preventive care would necessarily have developed the disease in question. Given this premise, the cost-effectiveness of preventive care depends on two factors: (1) the expected number of people for whom diseases could be prevented relative to the population at risk and (2) the cost of providing preventive treatments. Preventive treatments, such as *metformin*, can have low cost-effectiveness in cases where only a very small proportion of the population would have become ill without preventive care (Cohen et al. 2008). In this case, targeting patients who are at high risk of developing a disease and for whom there is a large differential effectiveness of preventive efforts could potentially improve the overall cost-effectiveness.

Previous research has assessed the cost-effectiveness of preventive care by comparing different preventive treatments or by comparing a preventive treatment against a placebo group (e.g., Breeze et al. 2017, Cohen et al. 2008, Zhou et al. 2020); but in these cases, the allocation rule was ex ante given and thus not the focus of the study. For instance, in the context of prediabetic patients, the cost-effectiveness of preventive care (here, lifestyle coaching) was studied by allocating the preventive treatment to patients aged 40–65, with a body mass index (BMI) higher than 35 and with blood glucose higher than 6 % (Breeze et al. 2017). However, the allocation rule used to determine when a prediabetic patient would receive the preventive care was devised by expert assessments and was not optimized through rigorous modeling. To fill this void, our work proposes a data-driven decision model.

2.2. Machine Learning in Health Management

Machine learning is widely used to predict health outcomes and to adapt the decision-making in health management accordingly. Given the extensive body of research, we provide only a short summary in the following paragraphs and point toward Keskinocak and Savva (2020) for a detailed overview. In health operations, predictions make widespread use of EHRs as input. EHRs have emerged as part of the recent digitalization of healthcare operations (KC et al. 2020); they encode a digital version of a patient's health trajectory, including a variety of risk variables, such as socio-demographics, body measurements (e.g., blood pressure), lab measurements, disease codes, and drug prescriptions. As a result, EHRs allow researchers to describe between-patient heterogeneity in health outcomes and, thereby, also enable a personalized view of a patient's individualized risk profile. Based on EHRs, health management can then predict different health outcomes. Examples of outcomes include the 30-day readmission risk as a performance indicator in hospital operations (e.g., Bardhan et al. 2015), mortality (e.g., Bjarnadottir et al. 2018), the onset of a disease (Choi et al. 2017), or complications (e.g., Schallmoser et al. 2023). In this work, we use machine learning for predicting a patient's risk of developing diabetes mellitus type II.

To predict the onset of diseases, research has used a variety of machine learning methods (e.g., Allam et al. 2021, Bertsimas et al. 2016, Zueger et al. 2022). Examples include lasso (Tibshirani 1996), ridge regression (Hoerl and Kennard 1970), random forests (Breiman 2001), neural networks (LeCun et al. 2015), and gradient boosted decision trees (Friedman 2001). Throughout this work, we use gradient boosted decision trees as our main model because this machine learning model has shown promising results in other management applications (e.g., Glaeser et al. 2019, Senoner et al. 2022). We later assess other machine learning methods as part of a sensitivity analysis.

2.3. Counterfactual Inference in Health Management

In health management, counterfactual inference is used to estimate the effects of different treatment options on health outcomes. Such estimates, in turn, inform the treatment choice, so that the treatment with the best expected effect on a patient's trajectory is chosen (Bica et al. 2021). The gold standard for estimating treatment effects is a randomized controlled trial, although it also has several shortcomings (Imbens and Rubin 2015). First, treatment effects may vary across patient populations, meaning that randomized controlled trials must be tailored to the patient population of interest. Second, randomized controlled trials are costly. Because of the immense upfront costs, performing such trials was not a viable course of action for our partnering IHO. Third, randomized controlled trials typically build on only two (or very few) treatment arms. As such, the estimated treatment effects typically are population-wide averages, with little personalization. Recent works (e.g., Bica et al. 2021) offer a remedy for these

shortcomings by using machine learning for estimating treatment effects from observational data, such as EHRs.

Previous works in counterfactual inference have pursued varying objectives (Guo et al. 2020). On the one hand are studies that make inferences at the group level and thus estimate average treatment effects (e.g., Frauen et al. 2023a, van der Laan and Rubin 2006, Shi et al. 2019). On the other hand are studies that estimate heterogeneous treatment effects based on a patient's risk score (Bica et al. 2021). Here, the assumption is that there is no global treatment effect of similar size for all patients but a differential treatment effectiveness—that is a heterogeneous treatment effect that varies across patients. As such, these studies may consider heterogeneous treatment effects related to, for instance, age, sex, or other risk scores. In the case of diabetes, preventive treatments such as *metformin* are known to have a differential effectiveness (Knowler et al. 2002). Thus, in our work, we use counterfactual inference to estimate heterogeneous treatment effects of *metformin* from EHRs.

We build on a recent method called causal forest (Wager and Athey 2018). The causal forest is a non-parametric method that extends the widely used random forest algorithm, by Breiman (2001), for estimating heterogeneous treatment effects. Causal forests inherit several favorable properties from random forests in that they are expressive, require little tuning, and have a low risk of overfitting, thus offering a robust performance in practice. Further, under common mathematical assumptions, causal forests have been shown to be pointwise consistent for the true treatment effect and thus lead to provably valid inferences (Wager and Athey 2018).

2.4. Resource Allocation in Health Management

Research on health management examines allocation problems related to various resources. For an overview, we refer to Dai and Tayur (2020) and Keskinocak and Savva (2020). Examples include beds (e.g., Helm et al. 2011), drugs (Khademi et al. 2015), aid (Jakubik and Feuerriegel 2022), and medical devices (Deo and Sohoni 2015). Other studies have used data-driven modeling for scheduling the admission of patients, so that available resources are used effectively but without exceeding surge capacities (e.g., May et al. 2011). Importantly, such problems generally optimize allocations according to patient needs (rather than focusing only on monetary aspects), so that providing (cost-)effective care is intended to improve patient health and not only to minimize costs. Analogously, we also focus on (cost-)effective care.

Allocation decisions have been extensively modeled in the context of secondary prevention, with the goal of optimizing screening or check-up policies. Examples include biopsy referral (Ayvaci et al. 2017), colonoscopy screening (Erenay et al. 2014), mammography (Cevik et al. 2018), and post-discharge monitoring (Helm et al. 2016, Liu et al. 2018). In this work, the main question is the optimal timing

for patient examinations (Kamalzadeh et al. 2021). Therefore, Markov decision processes or variants thereof typically are used in which (i) decisions for an individual patient are made and (ii) where states represent the patient's current health status. However, our problem setup for primary preventive care is different in both aspects: (i) decisions are made for a cross-sectional sample (i.e., determining which patients out of a cohort should be selected for treatment), and (ii) the future health trajectory should be used for decision-making, which requires machine learning models for prediction.

Ethical and legal frameworks around the world mandate whom to allocate treatments, including preventive treatments, given scarce resources (e. g., National Institute for Health and Care Excellence 2015). Accordingly, the core principle for health management is to focus on the best health benefit for patients, relative to the estimated costs, as opposed to prioritizing cost minimization alone. Hence, decisions in medicine are mainly benchmarked according to their cost-effectiveness (Gold et al. 1996): Decision models should optimize against the largest improvement in health outcomes, given budget constraints (typically set by policymakers or healthcare organizations).

3. Research Setting
3.1. The Case of Diabetes Prevention

Our research aims at allocation of preventive care to patients at risk of developing diabetes mellitus type II. Diabetes mellitus is a chronic condition and, although widespread, is often preventable. In the United States, 88 million adults, or 34.5 % of the adult population, are classified as prediabetic and thus at risk of developing diabetes if it is not prevented (U.S. Centers for Disease Control and Prevention 2020). According to the World Health Organization (2018a), diabetes mellitus not only is among the top 10 leading causes of death, but also can seriously imped the quality of life for patients in the long run. However, the risk of diabetes mellitus can be effectively mitigated with preventive treatments, including the drug *metformin* (Knowler et al. 2002). Once *metformin* has been prescribed and then taken for diabetes prevention, patients generally must remain on the medication for the rest of their lives.

In this research, we partnered with a national health insurer from Israel serving more than 1.1 million patients. In Israel, health insurers are vertically integrated healthcare organizations, which enables our partnering health insurer to have direct access to a wealth of medical data from EHRs, as well as to work collaboratively with healthcare providers to deliver care to patients. We later discuss the applicability of our model to vertically integrated healthcare organizations from other countries (see Section 6) and, for that reason, refer to our partner company simply as IHO.

Our IHO has set a goal of improving primary preventive care for patients at risk of developing diabetes mellitus type II to mitigate individuals' risk of onset. For this, we develop and evaluate a data-driven allocation of preventive care using *metformin*. Formally, the IHO is confronted by a

decision-making problem: which of its patients $i = 1, \ldots, N$ should be enrolled in preventive treatment $(t_i = 1)$ and which ones should not $(t_i = 0)$? If patient i is enrolled in preventive care, her risk of developing diabetes is reduced by a treatment effect, which we later estimate from historical EHRs. However, the IHO has limited the overall budget available for preventive care to $k \leq N$ treatments.

For primary prevention, the objective is to find a cost-effective allocation—that is, to maximize the number of prevented disease onsets, given a particular budget (World Health Organization 2018b). As with other research in health management (e. g., Ayvaci et al. 2012, Helm et al. 2011, 2016), the question is how to offer effective care at lower costs and thus to allocate resources for improving patient outcomes (as opposed to a goal of pure cost minimization, without considering health outcomes). Therefore, the objective in our work is to identify an allocation of preventive treatments to patients that maximizes the expected number of prevented onsets. To support the IHO in this task, we develop a data-driven decision model.

In addition to looking at the number of prevented disease onsets, we also report costs, thus allowing us to study cost-effectiveness. Note that we intentionally refrain from having only a cost minimization focus because of ethical and legal principles (National Institute for Health and Care Excellence 2015) and, as is common in medicine (Gold et al. 1996), we focus on cost-effectiveness. Depending on the success of preventive treatments, the IHO incurs different costs. A preventive treatment incurs an annual cost of C_{prevent} per patient. By enrolling all N patients in the preventive program, the IHO would experience a total cost of $N \times C_{\text{prevent}}$, which generally would exceed available financial resources. If the disease is not prevented, the IHO faces annual costs of treating a patient with diabetes, C_{diab}, which depends both on the age of the patient and on any concomitant diseases caused by the developed diabetes. The cost for treatment usually is substantially larger than the cost for prevention—that is, $C_{\text{diab}} \gg C_{\text{prevent}}$. However, not all prediabetic patients actually develop diabetes; hence, the cost C_{diab} occurs only for a subset of the overall study population.

3.2.   Electronic Health Records

Our work builds on an extensive longitudinal dataset. Specifically, our partnering IHO provided us with the EHRs of all patients classified as prediabetic and thus at risk of developing diabetes mellitus type II. Our sample comprises 89,191 patients and exhibits considerable between-patient heterogeneity in risk profiles. For purposes of this research, access to the EHRs was granted for the time period of 2003 through 2012.

For each patient, the EHRs comprise variables from the following categories: (i) socio-demographics; (ii) body measurements; (iii) lab tests; (iv) disease codes; and (v) drug prescriptions. Disease codes provide data on other co-occurring conditions and are encoded based on the International Classification

of Diseases (ICD) system of diagnostic codes. Table 1 shows examples for each category (i)–(v). These variables present potential risk factors that describe heterogeneity in disease onsets among patients, and we thus leverage them later for machine learning.

Table 1   **Variables in our dataset.**

Category	Number of variables	Examples
Socio-demographics	2	Age, sex
Body measurements	5	Height, weight, body mass index, systolic blood pressure, diastolic blood pressure
Lab tests	100	HbA1c, HDL-cholesterol, fasting glucose, . . .
Disease codes	100	Hypertensive diseases, disorders of metabolism, . . .
Drug prescriptions	100	Metformin, tritace, cardiloc, . . .
Overall (=input for our decision model)	307	

Patients were classified into (a) prediabetic or (b) diabetic as follows: (a) Prediabetic is the precursor stage before the onset of diabetes mellitus type II. Following established criteria, patients are considered prediabetic when the fasting glucose level is between 6.1 mmol/L and 6.9 mmol/L (World Health Organization 2006). The IHO used this inclusion criterion when compiling our patient data. (b) Diabetic corresponds to a fasting glucose level above 6.9 mmol/L (World Health Organization 2006). Later, during machine learning, our target variable denotes whether a patient has transitioned from prediabetic to diabetic ($=1$) or not ($=0$).

3.3. Summary Statistics

Our dataset comprises 89,191 prediabetic patients. Of these patients, 77,036 remain prediabetic throughout the study period, whereas 12,155 develop diabetes mellitus. Hence, only 13.62 % of patients transition from prediabetes to diabetes during the study period. Table 2 lists summary statistics for key variables for the first record of each patient in our dataset. In our study population, the average age is 47.07, and the average body mass index (BMI) is 29.05. These observations are in line with earlier findings from medical research, according to which (pre-)diabetes is especially prevalent in elderly patients and patients with high BMI (U.S. Centers for Disease Control and Prevention 2020).

Further differences arise when comparing patients who developed diabetes vs. patients who stayed prediabetic. We find that the average age is lower among patients who did not experience onset (46.84 years), compared to patients who developed diabetes (50.62 years). Similarly, the average BMI is larger for patients who developed diabetes (30.52 kg/m^2) compared to patients who stayed prediabetic (28.95 kg/m^2). Both observations are in line with medical research, which has identified age and BMI as important risk factors for the onset of diabetes (U.S. Centers for Disease Control and Prevention 2020).

10

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

Table 2   Summary statistics of example variables (for the first record of each patient in our dataset).

Variable	Overall		Without onset		With onset	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
Age (in years)	47.07	17.93	46.84	18.32	50.62	9.28
Sex (0=male; 1=female)	0.51	0.50	0.50	0.50	0.54	0.50
Systolic blood pressure (in mm Hg)	122.84	17.11	122.54	16.51	127.43	24.01
Diastolic blood pressure (in mm Hg)	77.26	12.78	77.16	12.93	78.80	10.12
Body mass index (in kg/m²)	29.05	5.64	28.95	5.62	30.52	5.73
HbA1c (in %)	5.56	0.36	5.51	0.34	5.87	0.34

SD = standard deviation

3.4.   Current Practice in Diabetes Prevention

In clinical practice, preventive care for diabetes is allocated as follows (Breeze et al. 2017): First, a health practitioner determines the patient's risk profile r_i and then calculates a risk scoring to assess her risk of developing diabetes. In practice, the risk score is computed using a simple charting tool (e. g., the Framingham diabetes risk score). Such charting tools have an obvious limitation in that only a few variables from the risk profile are considered; a holistic assessment with complete EHRs is absent. Afterward, the patient is enrolled in preventive care if the risk score exceeds a certain threshold ψ, i. e.,

$$t_i = \begin{cases} 1, & \text{if} \quad r_i \geq \psi, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

The threshold ψ is typically determined by experts in the field. To this end, we compute the threshold ψ for enrolling patients in preventive care, such that the number of patients enrolled is set to k per year—that is, the k patients with the highest risk scores are eligible for the preventive care program.

In our evaluation, the Framingham diabetes risk score typically is used to assess diabetes risk. The Framingham diabetes risk score was developed by the U. S. National Institutes of Health (Wilson et al. 2007) and represents the quasi-standard in clinical practice (Long and Fox 2016). The score is calculated using the charting tool shown in Table 3. Note that only a few risk factors are considered (i. e., fasting glucose level, BMI, cholesterol, parental history of diabetes mellitus, triglycerides, and blood pressure); all other variables from a patient's EHR are ignored. The failure to account for other variables thereby points toward opportunities for more effective risk assessment and prevention.

Table 3   Calculation of the Framingham diabetes risk score.

Risk factor	Score
If fasting glucose level between 100 and 126mg/dL	+10
If body mass index between 25.0 and 29.9	+2
If body mass index $\geq$ 30.0	+5
If cholesterol level $<$ 40 mg/dL in men or $<$ +50 mg/dL in women	+5
If parental history of diabetes mellitus	+3
If triglyceride level $\geq$ 150 mg/dL	+3
If blood pressure $\geq$ 130/85 mm Hg or receiving treatment	+2
Framingham diabetes risk score r_i	$\sum$

3.5.   Performance Metrics

We measure the operational performance of preventive care in such a way that the measure captures current practice at our IHO. Recall that preventive care aims to maximize the success of its preventive efforts (World Health Organization 2018b). In other words, given a certain budget, the expected number of disease onsets should be minimized. In medical decision-making, this measure is the quasi-standard in which outcome optimization is prioritized over cost minimization alone because of ethical and legal principles (National Institute for Health and Care Excellence 2015). In accordance with this standard, the expected number of prevented onsets represents our primary evaluation metric. Because health management also is interested in the economic value of implementing preventive care programs, we further report the projected cost savings but emphasize their secondary role.

We evaluate the performance of our decision model in a dynamic setting. Formally, we apply our decision model over a horizon of $l = 1, \ldots, L$ time steps (here, years 2003 through 2012). Each year allows access only to historical data to assign patients to preventive care. We denote that patient i was treated in year l by $t_{i,l} = 1$. The preventive treatment of patient i reduces the risk of developing diabetes in the future—that is, the treatment effect is $\gamma_{i,l}$. However, the outcomes of preventive treatment are not observable until months after the treatment. Let $y_{i,l+1} = 1$ denote the presence of diabetes for patient i in the following year, whereas $y_{i,l+1} = 0$ describes the absence of diabetes. To be realistic, we consider censoring analogously to how it occurs in practice. On the one hand, we account for patients who have died and exclude them from being targeted for preventive treatment. On the other hand, we do not discount patients where a follow-up was missed (and also do not exclude them from being targeted for a preventive treatment) because these patients can incur a cost in subsequent years. To this end, our performance metrics are as follows:

Prevented onsets. We evaluate our decision model by computing the expected number of prevented onsets, which directly corresponds to the objective of our decision model and of preventive care in practice. Hence, it is our prime metric of interest. The expected number of prevented onsets in the patient cohort amounts to

$$PreventedOnsets = \frac{1}{L} \sum_{l=1}^{L} \sum_{i=1}^{N} \gamma_{i,l} \, t_{i,l} \, y_{i,l+1}, \tag{2}$$

where $\gamma_{i,l} \, t_{i,l} \, y_{i,l+1}$ gives the expected reduction in the risk of a diabetes onset for patient i in year l, which is then averaged over the complete study horizon $1, \ldots, L$.

Cost savings. We calculate the expected cost savings of our allocation for our IHO as follows. The cost savings depend on the cost of enrolling patients in preventive care and the cost of treating diabetes. A preventive treatment (here: *metformin*) incurs an annual cost of $C_{\mathrm{prevent}} = \$1,380$ per patient (Gu et al. 2015). After a potential onset, the annual cost of treating diabetes depends on the

age of the patient, as well as additional comorbid conditions—that is, $C_{\mathrm{diab}} = C_0 + C_1 + \ldots + C_5$, where C_0 describes base costs that depend on the age of the patient, and $C_1, \ldots, C_5$ describe costs that can arise from comorbid conditions. In our case, we use $C_0 = \frac{\$15{,}000}{1+\exp(-age/10)}$ for base costs and consider five common comorbid conditions of diabetes: acute myocardial infarction ($C_1 = \$5{,}000$), intracerebral hemorrhage ($C_2 = \$5{,}000$), acquired hypothyroidism ($C_3 = \$5{,}000$), angina pectoris ($C_4 = \$15{,}000$), and heart failure ($C_5 = \$15{,}000$).

We then calculate the expected cost for patient i in year l when being treated using

$$C_{i,l}^{(\mathrm{t})} = y_{i,l+1} \left(1 - \gamma_{i,l}\right) C_{\mathrm{diab}} + C_{\mathrm{prevent}}, \tag{3}$$

and when not being treated using

$$C_{i,l}^{(\mathrm{nt})} = y_{i,l+1} \, C_{\mathrm{diab}}. \tag{4}$$

By summing over the patients and the study horizon, we compute the cost savings as

$$CostSavings = \underbrace{\sum_{l=1}^{L} \sum_{i=1}^{N} C_{i,l}^{(\mathrm{nt})}}_{\substack{\text{Cost for no} \\ \text{preventive treatments}}} - \underbrace{\sum_{l=1}^{L} \sum_{i=1}^{N} t_{i,l} \, C_{i,l}^{(\mathrm{t})} + (1 - t_{i,l}) \, C_{i,l}^{(\mathrm{nt})}}_{\substack{\text{Cost when allocating} \\ \text{patients to preventive program}}}. \tag{5}$$

Of note, some costs are ongoing and thus are incurred beyond the study horizon covered by our dataset. To account for this extension, we calculate the expected costs using an expected lifetime of 75 (Tachkov et al. 2020). For older patients, we take into account a minimum of three additional years of life, and for younger patients, we consider a maximum of ten years. In doing so, we reflect that the longer remaining lifetime of younger patients means that diabetes is likely to lead to larger overall costs for treatment.

4. Model Development
4.1. Problem Formulation

The objective of our partnering IHO is to dynamically allocate preventive care to patients so as to maximize the number of prevented cases of diabetes mellitus, given the available budget. Let $i = 1, \ldots, N_+$ refer to the customers of the IHO—that is, to the patients in our dataset. The plus symbol indicates that some of the patients have already received preventive treatments, whereas the majority has not been treated. Let the first $N < N_+$ customers denote patients who have not been treated yet and therefore potentially benefit from being enrolled in the preventive care program. Simply enrolling all patients in preventive care is not economically feasible because of budget constraints. Instead, in each year $l = 1, \ldots, L$, a decision must be made as to whether preventive care is warranted to patient i ($t_{i,l} = 1$) or is not warranted ($t_{i,l} = 0$).

For the purpose of decision-making, our IHO has access to additional information on patients' risk profiles through their EHR. EHR systems, similar to the one providing data for our evaluations, are

used worldwide, thus ensuring the broad applicability of our model. For simplicity, we refer to the variables in the EHR by $x_{i,l} \in \mathbb{R}^n$, $i = 1, \ldots, N_+$, $l = 1, \ldots, L$. These variables encode the heterogeneity among patients and are used as predictors in our data-driven approach. Specifically, we use the EHRs for identifying patients at risk of developing diabetes, which is relevant in clinical practice because only a small subset of prediabetic patients actually develop diabetes. Hence, by accurately identifying patients who will experience an onset, preventive treatments can be allocated to patients who would most benefit from them.

The objective is to maximize the expected effect of the preventive treatments. In other words, the expected number of disease onsets should be minimized within our time frame. As such, the decision of whether to enroll a specific patient in preventive care is driven only by the expected risk reduction of the preventive treatment for the patient (Paulweber et al. 2010). IHOs also might have to consider additional budget restrictions—not at the patient level, but at the population level. Hence, the decision problem is given by

$$\min_{\substack{t_{1,1}, \ldots, t_{1,L} \\ \cdots \\ t_{N,1}, \ldots, t_{N,L}}} \quad \sum_{l=1}^{L} \left(\sum_{i=1}^{N} t_{i,l} \, \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 1\right] + \sum_{i=1}^{N} (1 - t_{i,l}) \, \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 0\right] \right) \tag{6a}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} t_{i,l} \leq k, \quad \forall l \in \{1, \ldots, L\}, \tag{6b}$$

$$t_{i,l} \in \{0, 1\}, \quad \forall i \in \{1, \ldots, N\} \forall l \in \{1, \ldots, L\}. \tag{6c}$$

Equation (6a) encodes the objective whereby the expected number of disease onsets is minimized. Here, the expectation $\mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 1\right]$ denotes the probability of developing diabetes, conditional on the EHR of patient i in year l and in the presence of preventive treatment ($t_{i,l} = 1$). The expectation $\mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 0\right]$ denotes the probability of an onset, conditional on the EHR of patient i in absence of a preventive treatment ($t_{i,l} = 0$). Neither expectation is known, implying that the treatment effect must be estimated from data—for example, from electronic health records or, alternatively, by conducting randomized controlled trials. Equation (6b) introduces a budget constraint on the number of patients that can be enrolled in preventive care, thus allowing the treatment of k patients per year. Equation (6c) makes this problem an integer problem.

Our decision problem focuses on the expected number of diabetes onsets, similar to clinical practice; therefore, the cost of preventive treatments is not considered. Nevertheless, we can show that minimizing onsets also minimizes the overall cost of treatment (comprising both preventive treatments and treatments after a possible onset of the disease), given mild assumptions. This mathematical property is of direct relevance for health management because it underscores that improving the effectiveness of preventive care through better allocations also has financial value for healthcare organizations.

14

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

PROPOSITION 1 **(Link of prevented diseases to cost minimization)**.

Let $t_{1,l}, \ldots, t_{1,L}, \ldots, t_{N,1}, \ldots, t_{N,L}$ be a solution to the decision problem in Equation (6), let C_{prevent} denote the costs for prescribing preventive treatments, and let C_{diab} denote the costs for treating the disease after onset. Then, if

$$t_{i,l} \left(C_{\text{diab}} \, \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 1 \right] + C_{\text{prevent}} \right) \leq t_{i,l} \left(C_{\text{diab}} \, \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 0 \right] \right), \quad \forall i \in \{1, \ldots, N\} \forall l \in \{1, \ldots, L\}, \tag{7}$$

it follows that the solution $t_{1,l}, \ldots, t_{1,L}, \ldots, t_{N,1}, \ldots, t_{N,L}$ also minimizes costs for the health insurer.

Proof See Appendix A.

4.2. Proposed Decision Model

We develop a data-driven decision model for allocating preventive care—that is, a decision model that receives EHRs from patients as input and for which the output is a cost-effective allocation of preventive care resources (Figure 1). Our decision model uses a combination of causal inference, machine learning, and optimization, structured in three consecutive stages: (1) The first stage uses counterfactual inference to estimate the effectiveness of preventive treatments using EHRs (i.e., the heterogeneous treatment effect). Thus, we quantify how the personalized diabetes risk of a patient with a specific risk profile changes when a preventive treatment is prescribed. Here, we serve needs in practice, where the treatment effectiveness is often unknown but must be estimated from EHRs. Later, we also discuss an alternative approach that can be used when the treatment effect is known (e.g., from a randomized controlled trial). (2) The second stage estimates the personalized risk of a diabetes onset for an individual patient. Specifically, we apply machine learning to a patient's EHR to predict the probability of an onset. (3) The third stage uses the predicted risk together with the treatment effect to optimize the allocation of preventive care to patients under the given budget constraint.

Figure 1   **Decision model for data-driven allocations of preventive care.**
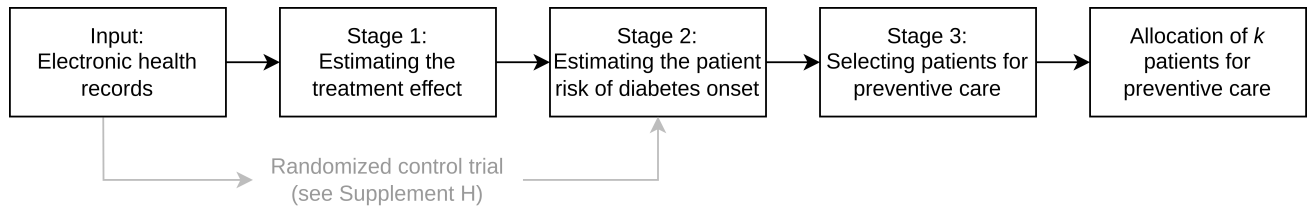
4.2.1. Stage 1: Estimating the Treatment Effect. In the first stage, we estimate the heterogeneous treatment effect of *metformin*—that is, the expected reduction in the probability of a diabetes onset when prescribing *metformin* to patients, given a specific patient risk profile. To account for heterogeneity in the treatment effectiveness across different patients, we introduce the following notation:

Let $\{1, \ldots, N_+\}$ be the set of patients in our sample, including patients who already received preventive treatment through *metformin*, as well as patients who have not received preventive treatments. Let $I = \{(i, l) : i \in \{1, \ldots, N_+\}, l \in \{1, \ldots, L\}\}$ be the set of 2-tuples for indexing patient-level observations throughout our study period.

We then estimate the treatment effect for each patient i in year l via

$$\gamma_{i,l} = \mathbb{E}\left[y \mid x_{i,l}, t_{i,l} = 0\right] - \mathbb{E}\left[y \mid x_{i,l}, t_{i,l} = 1\right]. \tag{8}$$

As a result, we allow the treatment effect to vary across patients who have different risk profiles. Building on the potential outcomes framework, it can further be shown that the treatment effects are identifiable from observational data (Rubin 2005).

To estimate the heterogeneous treatment effects, we use an established method: the causal forest (Wager and Athey 2018). The causal forest is a non-parametric method that extends the widely used random forest algorithm by Breiman (2001) for estimating heterogeneous treatment effects. This approach comes with many benefits. First, the causal forest is an ensemble method, which means it combines the predictions of multiple simpler models to improve its performance. As a result, the causal forest works well in settings where the treatment effect is subject to complex, nonlinear relationships. Second, under common mathematical assumptions, the causal forest has been shown to be pointwise consistent for the true treatment effect and thus leads to provably valid inferences. Third, the causal forest is capable of handling high-dimensional and mixed data types with little risk of overfitting, making it a robust choice that requires minimal preprocessing. Mathematically, one of these characteristics is its honesty property, where the trees in the causal forest are grown using one subsample of the dataset, while the predictions at the leaves of the tree are estimated using a different subsample (Wager and Athey 2018). By ensuring that each tree is only fit to a portion of the data, the honesty property helps to mitigate the risk of overfitting and to produce accurate and reliable treatment effect estimates. We report details on our implementation in Section G.

As part of our robustness checks, we perform a causal sensitivity analysis (Cinelli and Hazlett 2020) in which we study the effect of potential unobserved confounding variables on the estimated treatment effects. We find that, for unobserved confounding variables that are even strongly associated with our treatment, the treatment effects remain robust (see Supplement D). Furthermore, we study how errors in the treatment effect estimation influence the overall performance of our decision model, and we find that the results remain largely robust. Finally, we provide an extension to settings where the treatment effect is not estimated from historical data but is given ex ante (e.g., where treatment effect is obtained through a randomized controlled trial instead), and we show how our decision model is applicable to such cases (see Supplement H).

4.2.2.   Stage 2: Estimating the Patient Risk of Diabetes Onset. In the second stage, we rely on machine learning to estimate the heterogeneous risk of diabetes onset (in the absence of preventive treatment). For this purpose, we first filter out patients who already received preventive treatment. Let the first $N < N_+$ customers be patients who have not already been enrolled in a preventive care program. We then train a machine learning model h_0^θ with parameters $\theta \in \Theta$ as follows. The machine learning model uses EHRs, $x_{i,l}$, to predict the risk of the individuals' developing the disease. Formally, this is given by $\mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 0\right]$, $i \in \{1, \ldots, N\}$, $l \in \{1, \ldots, L\}$. Hence, we have $h_0^\theta(x_{i,l}) = 1$ if there is an onset in the following year, and $h_0^\theta(x_{i,l}) = 0$ if the condition for patient i is absent in the following year. Given h_0^θ and the estimated treatment effect $\gamma_{i,l}$ from stage 1, we can then compute the probability of onset, conditional on a preventive treatment. Then, the risk of a diabetes onset under preventive treatment amounts to $(1 - \gamma_{i,l})\, h_0^\theta(x_{i,l})$.

Mathematically, h_0^θ is estimated as follows. Let $\theta \in \Theta$ refer to a set of parameters in the machine learning model h_0^θ. For estimation, we minimize the empirical risk; that is,

$$h_0^* = \operatorname*{arg\,min}_{\theta \in \Theta} \sum_{l=1}^{L} \sum_{i=1}^{N} \mathcal{L}(h_0^\theta(x_{i,l}), y_{i,l+1}), \tag{9}$$

where $\mathcal{L}$ denotes a loss function that measures the error of h_0^θ in forecasting the onset of a disease. We further calibrate the output probabilities of h_0^θ using Platt scaling (see Supplement F).

We later report results from our decision model using gradient-boosted decision trees. Gradient-boosted decision trees (Friedman 2001) belong to the category of ensemble methods, which are known to perform well on complex datasets and have been used in other operational applications (e. g., Glaeser et al. 2019). Nevertheless, we also perform an extensive series of robustness checks, repeating the analysis with other machine learning methods (e. g., lasso, ridge regression, random forest, and deep neural network). For all machine learning models, hyperparameters were tuned using 10-fold cross-validation. The tuning procedure is reported in Supplement G.

4.2.3.   Stage 3: Selecting Patients for Preventive Care. To accommodate budgetary constraints, we model the maximum number of preventive treatments as k patients per year. Our objective is to identify a cost-effective allocation that minimizes the expected number of onsets of the disease. Hence, we formulate our optimization problem as

$$\min_{\substack{t_{1,1},\ldots,t_{1,L} \\ \cdots \\ t_{N,1},\ldots,t_{N,L}}} \quad \sum_{l=1}^{L} \left(\sum_{i=1}^{N} t_{i,l}\, (1 - \gamma_{i,l})\, h_0^*(x_{i,l}) + \sum_{i=1}^{N} (1 - t_{i,l})\, h_0^*(x_{i,l}) \right) \tag{10a}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} t_{i,l} \leq k, \quad \forall l \in \{1, \ldots, L\}, \tag{10b}$$

$$t_{i,l} \in \{0, 1\}, \quad \forall i \in \{1, \ldots, N\} \forall l \in \{1, \ldots, L\}. \tag{10c}$$

This optimization problem extends Equations (6a) to (6c), in that we replace the (unknown) expected values with their corresponding predictions from the machine learning methods.

In the case of *metformin*, the solution to the previous optimization problem is the subset of k patients per year for which the estimated reduction in expected onset is higher than the reduction for patients not in the subset. We show this property in the following proposition.

PROPOSITION 2 **(Optimal allocation).** *Let h_0^* be an oracle for the risk of an onset, and let $\gamma_{i,l}$ be the true heterogeneous treatment effect of patient i in year l. The allocation that minimizes the expected number of onsets in the population $\{1, \ldots, N\}$ provides preventive treatments in year l to patients $\{1, \ldots, k\}$, where*

$$(1 - \gamma_{i,l}) \, h_0^*(x_{i,l}) \geq (1 - \gamma_{j,l}) \, h_0^*(x_{j,l}), \qquad \forall i \in \{1, \ldots, k\} \, \forall j \in \{k+1, \ldots, N\}. \tag{11}$$

Proof   See Supplement B.

In our implementation, we follow best practices and assess uncertainty in our model using a bootstrap procedure (Hastie et al. 2009). Specifically, we use 100 samples with replacement and then report both the mean and the standard deviation.

5.   Empirical Results
5.1.   Performance of Data-Driven Allocation

In this section, we evaluate our decision model and compare its performance against the clinical baseline, which we use to mimic healthcare practitioners who base their decision-making on the Framingham diabetes risk score (see Section 3.4). In addition, we report a naïve baseline, in which the allocation of preventive treatments is not optimized but, instead, a subset of k patients from the population is randomly chosen for each year. By comparing the naïve baseline with our decision model, we can directly quantify the relative gain from stratifying the allocation to high-risk patients through a data-driven approach. The results are in Table 4.

The naïve baseline (i.e., random allocation) is estimated to prevent 489 onsets when treating k=10,000 patients per year. For the same budget allowing for k=10,000 preventive treatments, the clinical baseline is estimated to prevent 701 onsets. Meanwhile, our decision model would prevent 882 onsets—an improvement of more than 25 percent. The results for other budgets lead to similar conclusions. In sum, our decision model is considerably more effective than current practice in disease prevention.

The use of our decision model offers considerable cost savings. A random allocation of preventive treatments leads to cost savings of \$6.319 million at k=1,000 and \$11.988 million at k=10,000. This

18

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

outcome shows that preventive care can be economical, in that the cost for treating a diabetes onset is substantially larger than the cost for preventive treatments. The clinical baseline improves cost savings further: to \$7.340 million at $k=1,000$ and to \$15.350 million at $k=10,000$. In comparison, our decision model demonstrates consistent cost savings over both the naïve baseline and the clinical baseline across all budgets k. For $k=1,000$ and $k=10,000$, our IHO could generate additional savings, beyond that of the clinical baseline, of \$0.710 million and \$2.987 million, respectively. As a result, the cost savings in our decision model are substantially larger than those of the clinical baseline and, based on a comparison with t-tests, are statistically significant at common significance thresholds.

Table 4   **Performance of our data-driven decision model.**

	$k=1,000$		$k=5,000$		$k=10,000$	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Naïve baseline	53.438	6.319	250.422	8.889	489.331	11.988
	(0.330)	(0.235)	(1.698)	(0.258)	(2.162)	(0.249)
Clinical baseline	134.798	7.340	446.948	11.996	701.942	15.350
	(1.025)	(0.245)	(3.373)	(0.258)	(3.837)	(0.265)
Our decision model	162.913	8.050	567.710	13.979	882.004	18.337
	(0.902)	(0.239)	(2.612)	(0.253)	(3.934)	(0.284)

Stated: mean performance (standard deviation in parentheses)

Note. Performance metrics for allocating preventive care given a varying budget for enrolling k patients per year into preventive treatments. Cost savings over no preventive care allocation are reported in USD millions.

5.2. Comparison of Patient Characteristics

In this section, we provide a quantitative assessment of the difference between the allocations from the clinical baseline and from our decision model. For this work, we report descriptive statistics for patients who were enrolled in preventive care (see Table 5). The clinical baseline based on the Framingham diabetes risk score allocates preventive care to patients who, on average, are 41.36 years old, whereas our decision model allocates preventive care to patients who are, on average, 53.60 years old and thus substantially older. This increase is in line with medical research (Knowler et al. 2002), which shows that a larger treatment effect from *metformin* is observed in people above age 45. Here, we note that the clinical baseline based on the Framingham diabetes risk score does not consider age to be a predictor of diabetes onset (Wilson et al. 2007), and for this reason, following current practice does not explicitly stratify the allocation of preventive care to older age groups. Furthermore, the clinical baseline based on the Framingham diabetes risk score selects patients with an average BMI of 31.31, while our decision model gives preference to patients with a larger BMI (average: 34.55). We discussed our observations with medical professionals specializing in diabetes care, who explained that the stratification from our decision model is beneficial because, in their experience, *metformin* leads to reduction in body weight;

they noted that such preventive efforts are especially effective when offered to overweight patients. In terms of gender, both allocations exhibit a similar distribution among men and women.

Table 5   **Descriptive statistics of patients enrolled in preventive care for different allocations.**

| Variable	Overall population		Clinical baseline		Our decision model	
	Mean	SD	Mean	SD	Mean	SD
Age (in years)	47.07	17.93	41.36	7.69	53.60	27.03
Sex (0=male; 1=female)	0.51	0.50	0.53	0.50	0.55	0.50
Systolic blood pressure (mm Hg)	122.84	17.11	124.86	14.19	125.65	17.13
Diastolic blood pressure (mm Hg)	77.26	12.78	78.72	9.76	78.98	11.90
Body mass index (kg/m^2)	29.05	5.64	31.31	6.16	34.55	6.10
HbA1c (in %)	5.56	0.36	5.74	0.28	5.83	0.34

SD = standard deviation; k=5,000

5.3.   Cost-Effectiveness Analysis

In our cost-effectiveness analysis (similar to Ayvaci et al. 2012), we estimate the number of prevented disease onsets and cost savings across varying budgets—that is, for a varying number of patients being enrolled in preventive care.

With increasing budgets, the relative proportion of prevented disease onsets also increases (see Figure 2). However, the improvement is more pronounced for our decision model, compared to both the naïve baseline and the clinical baseline. Hence, for the same budget increase, our decision model is estimated to be more effective in preventing disease onsets. For instance, when enrolling 10 % of the prediabetic population in preventive care, we estimate that the clinical baseline can prevent 2.21 % of diabetes onsets. Using our decision model, this proportion increases to 2.89 %.

Figure 2   **Prevented onsets when providing preventive care under different budget constraints.**
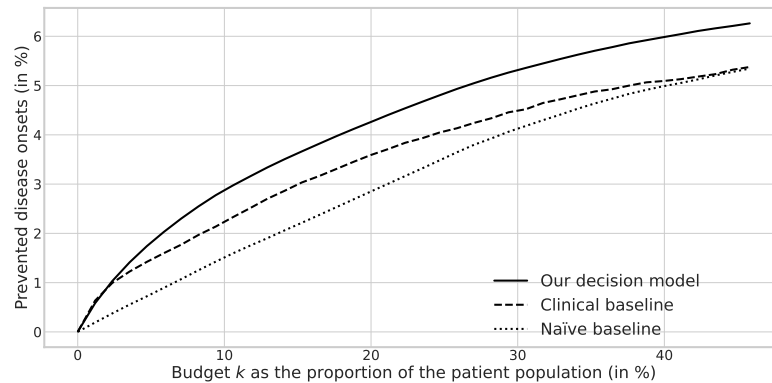
The cost savings over no preventive care are shown in Figure 3. As expected, the naïve baseline and the clinical baseline based on the Framingham diabetes risk score achieve lower cost savings than our decision model, confirming again the superiority of our decision model. This finding is consistent across different budgets (i. e., different numbers of patients enrolled in preventive treatments). Our decision

20

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

model leads to annual cost savings of up to \$20.31 per prediabetic patient. The cost savings per patient reach a maximum when $47.58\,\%$ of the population is enrolled in preventive treatments.

Figure 3   Cost savings when providing preventive care under different budget constraints.
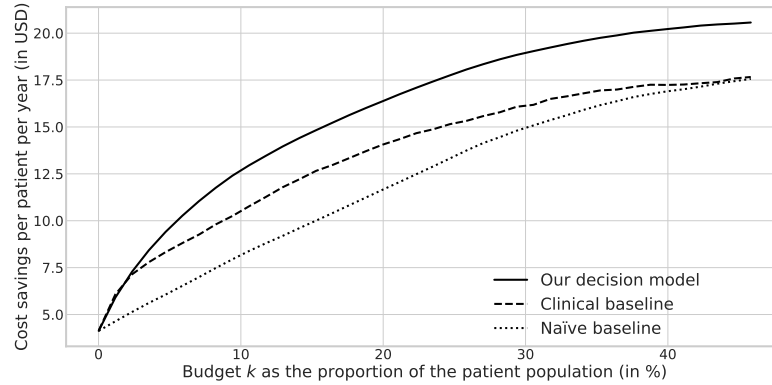
5.4.   Application to a U.S.-Wide Population

For better comparability, we now repeat the previous analysis for a population-wide setting, where we estimate the results of applying our decision model to the entire U.S. population. The U.S. population includes about 300 million adults, of whom 88 million are considered prediabetic. As previously noted, we assume that costs for preventive treatment are \$1,380 per patient and that model costs are dependent on age and potential comorbid conditions of a patient (see Section 3.5). We further assume the same machine learning performance in identifying high-risk patients as in the previous analyses.

On the basis of this setting, we obtain the following findings: Using the clinical baseline, allocating preventive treatments to $10\,\%$ of patients (equal to 8.8 million people) leads to cost savings of \$913 million annually. This outcome corresponds to annual cost savings of \$10.38 per prediabetic patient. In comparison, our decision model offers larger benefits. When providing $10\,\%$ of patients with preventive treatments, our decision model allows for annual cost savings of \$1.1 billion, equivalent to \$12.50 per prediabetic patient. Because the prevalence of diabetes in the United States is increasing (Chen et al. 2011), the cost savings are likely to continue to grow in the coming years.

5.5.   Robustness Checks

5.5.1.   Sensitivity analyses. We performed an extensive series of additional empirical analyses that add to the robustness of our findings: (1) We tested our decision model with varying machine learning models (e.g., random forest, deep neural network). Details are in Supplement C.1. We find that all machine learning models perform well, and furthermore, the results favor our choice of gradient-boosted decision trees. (2) We tested alternative risk scores (see Supplement C.2). Here, we find that our proposed decision model consistently leads to a larger number of prevented diseases and larger cost

savings. (3) We evaluated the robustness of the decision model to estimation errors in the treatment effect (see Supplement C.3). We again find that our decision model shows robust results—in particular, when there is a larger number of patients in the treatment group. (4) We demonstrate the applicability of our approach to personalized coaching as another form of preventive care (Supplement H).

5.5.2. Explainability. We further build on explanation methods to analyze how our machine learning model arrives at predictions of diabetes onsets. To identify the most important predictors in our prediction model, we calculated SHapley Additive exPlanations (SHAP) values (Lundberg and Lee 2017). We find that the variables of age, BMI, and HbA1c are the most important predictors for our machine learning model when assessing the risk of a diabetes onset. Larger values for these variables result in the machine learning model's calculating a larger probability of diabetes onset. We discussed the results with medical professionals, who confirmed that our findings not only match their own experience but also are in line with medical research. Details are in Supplement E.

5.5.3. Alternative Model Specifications. We further benchmark different variations of our decision model to confirm the effectiveness of our model specification:

1. *Main model.* We report the decision model from the main analysis. Here, the machine learning stage uses 307 input variables (see Table 1), and the optimization stage stratifies the allocation on the basis of risk reduction.

2. *Sparse model I/II.* For Sparse Model I, we vary the input as follows: The model uses only a subset of input variables—namely, those from the Framingham diabetes risk score (see Table 3). For Sparse Model II, we also use the risk factors of age and the current glycated hemoglobin level (HbA1c). This addition allows us to quantify the performance gain from using modern, high-dimensional EHRs.

3. *Linear model.* Here, we vary the machine learning stage by using only a linear model (with lasso regularization). By comparing this model against our main decision model (with gradient-boosted decision trees), we can assess the gains from modeling complex, nonlinear relationships (e. g., from comorbidities and other co-occurring conditions).

4. *Risk-stratified model.* For the risk-stratified model, we vary the optimization stage, stratifying the allocation only by patient risk. That is, we select the patients for preventive care based on the highest risk of an onset, without considering heterogeneous treatment effects.

Table 6 presents the results for the five different model specifications. (1) We find that the decision model from the main analysis performs best. (2) We find that using high-dimensional patient data from modern EHRs is beneficial, but only to a relatively small degree. Replacing the 307 input variables with the variables from the Framingham diabetes risk score, together with age and HbA1c, can still

lead to large improvements over the clinical baseline. In particular, the sparse model II is ranked second overall, even though it has only seven input variables. Hence, the use of advanced machine learning, together with an optimization stage that accounts for heterogeneous treatment effects of preventive care, is more important than having access to rich EHR data. (3) We find that nonlinear machine learning outperforms a linear model. The finding confirms the value of modeling complex, nonlinear relationships through nonlinear machine learning models. (4) We obtain large benefits from including heterogeneous treatment effects in the objective. Hence, the optimization stage is more effective when allocation decisions take into account risk reduction, rather than only patient risk.

Table 6   Results for different model specifications.

	k=1,000		k=5,000		k=10,000	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Clinical baseline	134.798	7.340	446.948	11.996	701.942	15.350
	(1.025)	(0.245)	(3.373)	(0.258)	(3.837)	(0.265)
Main model	162.913	8.050	567.710	13.979	882.004	18.337
	(0.902)	(0.239)	(2.612)	(0.253)	(3.934)	(0.284)
Sparse model I	137.209	6.745	466.106	11.364	728.151	15.041
	(0.757)	(0.195)	(2.152)	(0.210)	(3.326)	(0.234)
Sparse model II	159.148	7.850	549.777	13.492	855.743	17.752
	(0.880)	(0.231)	(2.532)	(0.246)	(3.848)	(0.276)
Linear model	152.124	7.489	520.490	12.726	811.841	16.800
	(0.840)	(0.218)	(2.401)	(0.234)	(3.684)	(0.261)
Risk-only model	145.354	7.148	494.492	12.063	772.251	15.958
	(0.802)	(0.207)	(2.282)	(0.222)	(3.523)	(0.249)

Stated: mean performance (standard deviation in parentheses)

Note. Performance metrics for allocating preventive care given a varying budget for enrolling k patients per year into preventive treatments. Cost savings over no preventive care allocation are reported in USD millions.

6. Discussion
6.1. Implications for Research

Our novel decision model for data-driven allocation of preventive care integrates counterfactual inference, machine learning, and optimization. Specifically, the model allocates a given budget toward primary preventive care for patients at risk of developing a disease, so as to maximize the expected number of patients whose disease onset is prevented. We evaluate our decision model in a dynamic setting using a large, longitudinal dataset of 89,191 patients at risk of developing diabetes mellitus type II, thereby confirming its effectiveness.

We study the operational value of incorporating high-dimensional health data, machine learning, and optimization toward preventive treatment allocation decision-making. Through our experiments using a large, longitudinal dataset, we examined the performance gap between our proposed decision model and alternative model specifications (see Section 5.5.3). Overall, the largest performance gain in prevented diseases and cost savings comes from targeting patients based on risk reduction rather than only on risk. Hence, a crucial implication of our work is the importance of a rigorous modeling framework for decision-making. Also crucial is the use of advanced machine learning to accommodate

complex, nonlinear relationships. A comparatively smaller gain comes for the use of high-dimensional EHRs. In fact, a decision model that uses machine learning and targets patients on the basis of risk reduction, but that has access only to a small subset of patient variables (sparse model II), can even achieve a performance that is close to that of our main model. Our findings imply that using data-driven decision models to inform resource allocation in healthcare may be practical, even in cases where relatively little information is available about a patient.

Proactively reducing individuals' risk of developing preventable diseases is a top priority in healthcare not only because it can improve patients' quality of life but also because it reduces healthcare costs. Our analysis shows that a large number of onsets can be prevented by using a data-driven decision model. In practice, allocation of patients to preventive care programs is often solely based on the risk of disease onset for the patients—that is, patients with the highest risk of onset are assigned to preventive treatments. However, prior medical literature has shown that preventive treatments can have a differential effectiveness (Knowler et al. 2002). By exploiting this differential effectiveness, our method allocates preventive treatments in a more cost-effective manner, including prioritizing patients who have a lower risk of onset but who can benefit from preventive care more than patients who have a higher risk of onset. Similar observations have been made for targeting customers in marketing, where interventions can benefit when they consider not only the risk that a customer is about to churn, but also the incremental effectiveness of marketing interventions (Ascarza 2018).

In our decision model, we also address the problem of unobserved confounders using a powerful statistical framework (Cinelli and Hazlett 2020). Unobserved confounders simultaneously influence both the probability of being assigned to preventive treatment and the likelihood of developing diabetes. Thus, if unobserved confounders are not accounted for, the estimates of the treatment effect can be biased (Frauen et al. 2023b). Such confounders are common in medicine and can include such variables as patients' socioeconomic status, healthcare coverage, and race. Often, these variables either cannot be directly observed or cannot be collected and used because of anti-discrimination laws. In light of this limitation, statistical frameworks that can address the potential effects of unobserved confounders can offer significant practical value in operations management.

6.2. Implications for Health Management

Healthcare managers have had to confront rising costs for care in recent years (KC et al. 2020). To counteract these costs, regulatory initiatives, such as the Affordable Care Act (§4001(f)(3)), dictate a greater focus on disease prevention (Koh and Sebelius 2010). Health management research has long studied the efficient provision of care for patients *after* they develop a medical condition (e.g., Garg et al. 2020, Zhang et al. 2016); in contrast, our work charts a different path, where care is allocated

prior to any disease onset, with the goal of preventing it. To support health management in this task, we develop a novel data-driven decision model.

Current practice in allocating preventive care is based on risk scores and has several significant shortcomings. First, traditional risk scores (e. g., charting tools) are limited to a small subset of relevant patient variables that are commonly available in contemporary EHRs. Here, leveraging a few key variables from modern EHRs can offer better predictive power in identifying patients at risk of developing diabetes. Second, traditional risk scores also are limited to simplified linear models thus neglecting complex, nonlinear relationships between risk factors. Yet, such nonlinear relationships are particularly common in the presence of comorbidities and other co-occurring conditions. Third, traditional risk scores allocate preventive care according to a patient's risk level but without considering the heterogeneous treatment effects of preventive treatments across patients—that is, the differential effectiveness in reducing the risk of onset for a patient. Our work's contributions address these three key limitations.

6.3. Practical Considerations

Our decision model can be readily developed and integrated into existing health management practices. Specifically, the approach we develop here is a general-purpose framework that can produce cost-effective intervention programs by healthcare systems that have access to EHRs; our approach can be applied to maximize prevention of any preventable disease for which preventive treatment exists, such as cardiovascular diseases or respiratory diseases, and for a given preventive care budget. As such, our approach is neither limited to a particular preventable disease nor does it rely on having access to a specific kind of information. In this work, we present empirical evaluations when using real healthcare data commonly recorded in practice by standard EHR systems worldwide. Therefore, our decision model is applicable to other IHOs in which similar standard EHR systems are used.

In recent years, there has been a transition in healthcare from reactive to proactive care, underscoring the importance of preventive practices. Proactive care is a comprehensive approach that focuses on identifying and mitigating risk factors to prevent the onset of diseases, in contrast to only treating diseases and conditions once they arise. By proactively reducing individuals' risks of developing preventable diseases, health management can improve patient outcomes and reduce the long-term cost and resource burden on the healthcare system (Zhou et al. 2020). This approach aligns with the objectives of IHOs in many countries, including several European countries, Israel, Singapore, and others. (World Health Organization 2016). In the U. S., some healthcare systems have goals that are consistent with cost-effective interventions to benefit the long-term well-being of patients; they include health maintenance organizations (e.g., the Kaiser group, with over 12 million members (e.g., Kaiser Permanente

2021)) and Accountable Care Organizations, which merge patient data from different providers, with the goal of promoting cost-effective care. Generally, IHOs aim to improve coordination and efficiency among healthcare services and typically prioritize preventive care and patients' long-term well-being more broadly. By having access to comprehensive patient information from different health providers, IHOs can make informed decisions to optimize the allocation of resources for preventive care. Our decision model is directly applicable in such contexts and thus facilitates proactive care in two ways: (1) by providing a data-driven, cost-effective allocation of limited preventive care resources, and (2) by supporting health managers with the goal of improving patient outcomes and mitigating the burden of preventable debilitating diseases. Worldwide, evidence on the value of care and the use of data-driven care have led to an increasing trend toward integration in the healthcare sector (e.g., World Health Organization 2016, Cuesta et al. 2019). In contrast, in countries where IHOs are not available, our decision model can motivate and promote collaboration among providers and insurers to enable cost-effective preventive care.

Operationally, the frequency of risk assessment and subsequent (re)allocation of a preventive treatment to patients ought to reflect the particular context. For instance, the onset risk of many preventable chronic diseases, including metabolic diseases such as diabetes mellitus type II, can vary over time, and a periodic estimation of risks and allocation of treatments, as we have done in this work, would be prudent. We have demonstrated how annual reallocation of preventive treatments can meaningfully reduce the number of patients affected by diabetes mellitus. In other contexts, the updating frequency ought to reflect the rate at which patients' risk is likely to change. As we have demonstrated here, our decision model can be applied to yield updated risk updates and to reallocate the preventive treatments accordingly.

Primary prevention considers all patients and prioritizes offering intervention treatments to patients based on their respective expected risk reduction. Thus, preventive care can yield the greatest benefits when it is not limited to considering only a subset of patients who initiate a physician's visit; instead, it could be proactively offered to patients for whom preventive care can be most effective, including patients who may be unaware of their own risk and thus do not seek treatment. In particular, given that nine in ten U.S. adults are unaware of their high-risk status (Geiss et al. 2010), proactive efforts are paramount for effective diabetes prevention.

Implementing a decision model for the allocation of preventive care is subject to several practice challenges. One challenge is to ensure the availability of EHRs, which is a crucial requirement for identifying patients at risk of disease onset and for estimating heterogeneous treatment effects. Another challenge is the quality of EHR data, especially with respect to completeness and accuracy, both of

26   **Author:** *Data-Driven Allocation of Preventive Care*

Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

which can affect the performance of the decision model. Finally, there may be regulatory challenges in implementing the decision model, such as privacy and data protection regulations that restrict access to patient data. As in any data-driven approach, establishing that the use of information to allocate disease prevention is consistent with local privacy and regulatory frameworks is important. Future work may also extend our analysis through the use of algorithmic fairness (De-Arteaga et al. 2022).

6.4. Concluding Remarks

In an effort to reduce costs and improve quality of life, health management places an increasing focus on disease prevention. However, healthcare organizations lack the decision support necessary to facilitate effective allocation of preventive treatments to patients, given that population-wide access of preventive care is often prohibitively costly. To aid health management in this task, we have developed in this work a data-driven decision model for the cost-effective allocation of preventive care.

Acknowledgments

References

Abbasi A, Sahlqvist AS, Lotta L, Brosnan JM, Vollenweider P, Giabbanelli P, Nunez DJ, et al. (2016) A systematic review of biomarkers and risk of incident type 2 diabetes: An overview of epidemiological, prediction and aetiological research literature. *PLOS ONE* 11(10):e0163721.

Allam A, Feuerriegel S, Rebhan M, Krauthammer M (2021) Analyzing patient trajectories with artificial intelligence. *Journal of Medical Internet Research* 23(12):e29812.

American Diabetes Association (2022) 2. Classification and diagnosis of diabetes: Standards of medical care in diabetes. *Diabetes Care* 45(Supplement1):S17–S38.

Ascarza E (2018) Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research* 55(1):80–98.

Ayvaci MUS, Ahsen ME, Raghunathan S, Gharibi Z (2017) Timing the use of breast cancer risk information in biopsy decision-making. *Production and Operations Management* 26(7):1333–1358.

Ayvaci MUS, Alagoz O, Burnside ES (2012) The effect of budgetary restrictions on breast cancer diagnostic decisions. *Manufacturing & Service Operations Management* 14(4):600–617.

Azelton KR, Crowley AP, Vence N, Underwood K, Morris G, Kelly J, Landry MJ (2021) Digital health coaching for type 2 diabetes: Randomized controlled trial of healthy at home. *Frontiers in Digital Health* 3.

Bardhan I, Oh JH, Zheng Z, Kirksey K (2015) Predictive analytics for readmission of patients with congestive heart failure. *Information Systems Research* 26(1):19–39.

Bertsimas D, O'Hair A, Relyea S, Silberholz J (2016) An analytics approach to designing combination chemotherapy regimens for cancer. *Management Science* 62(5):1511–1531.

Bica I, Alaa AM, Lambert C, van der Schaar M (2021) From real-world patient data to individualized treatment effects using machine learning: Current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics* 109(1):87–100.

Bjarnadottir M, Anderson D, Zia L, Rhoads K (2018) Predicting colorectal cancer mortality: Models to facilitate patient-physician conversations and inform operational decision making. *Production and Operations Management* 15(4):2162–2183.

Bolnick HJ, Bui AL, Bulchis A, Chen C, Chapin A, Lomsadze L, Mokdad AH, et al. (2020) Health-care spending attributable to modifiable risk factors in the USA: An economic attribution analysis. *The Lancet Public Health* 5(10):525–535.

Breeze PR, Thomas C, Squires H, Brennan A, Greaves C, Diggle PJ, Brunner E, et al. (2017) The impact of type 2 diabetes prevention programmes based on risk-identification and lifestyle intervention intensity strategies: A cost-effectiveness analysis. *Diabetic Medicine* 34(5):632–640.

Breiman L (2001) Random forests. *Machine Learning* 45(1):5–32.

Cevik M, Ayer T, Alagoz O, Sprague BL (2018) Analysis of mammography screening policies under resource constraints. *Production and Operations Management* 27(5):949–972.

Chaker L, Ligthart S, Korevaar TI, Hofman A, Franco OH, Peeters RP, Dehghan A (2016) Thyroid function and risk of type 2 diabetes: A population-based prospective cohort study. *BMC Medicine* 14(150).

Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP (2002) SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* 16:321–357.

Chen L, Magliano DJ, Zimmet PZ (2011) The worldwide epidemiology of type 2 diabetes mellitus–present and future perspectives. *Nature Reviews Endocrinology* 8(4):228–236.

Choi E, Schuetz A, Stewart WF, Sun J (2017) Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association* 24(2):361–370.

Cinelli C, Hazlett C (2020) Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82(1):39–67.

Cohen JT, Neumann PJ, Weinstein MC (2008) Does preventive care save money? Health economics and the presidential candidates. *The New England Journal of Medicine* 358(7):661–663.

Cuesta JI, Noton C, Vatter B (2019) Vertical integration between hospitals and insurers. *Available at SSRN 3309218* .

Dai T, Tayur S (2020) Healthcare operations management: A snapshot of emerging research. *Manufacturing & Service Operations Management* 22(5):869–887.

De-Arteaga M, Feuerriegel S, Saar-Tschansky M (2022) Algorithmic fairness in business analytics: Directions for research and practice. *Production and Operations Management* 31(10):3749–3770.

Deo S, Sohoni M (2015) Optimal decentralization of early infant diagnosis of HIV in resource-limited settings. *Manufacturing & Service Operations Management* 17(2):191–207.

Erenay FS, Alagoz O, Said A (2014) Optimizing colonoscopy screening for colorectal cancer prevention and surveillance. *Manufacturing & Service Operations Management* 16(3):381–400.

Frauen D, Hatt T, Melnychuk V, Feuerriegel S (2023a) Estimating average causal effects from patient trajectories. *Proceedings of the AAAI Conference on Artificial Intelligence*, 7586–7594.

Frauen D, Melnychuk V, Feuerriegel S (2023b) Sharp bounds for generalized causal sensitivity analysis. *arXiv 2305.16988* .

Friedman JH (2001) Greedy function approximation: a gradient boosting machine. *Annals of Statistics* 1189–1232.

Garg T, Bhardwaj M, Deo S (2020) Role of community health workers in improving cost efficiency in an active case finding tuberculosis programme: An operational research study from rural Bihar, India. *BMJ Open* 10(10):e036625.

Geiss LS, James C, Gregg EW, Albright A, Williamson DF, Cowie CC (2010) Diabetes risk reduction behaviors among U.S. adults with prediabetes. *American Journal of Preventive Medicine* 38(4):403–409.

Glaeser CK, Fisher M, Su X (2019) Optimal retail location: Empirical methodology and application to practice. *Manufacturing & Service Operations Management* 21(1):86–102.

Goetzel RZ (2009) Do prevention or treatment services save money? The wrong debate. *Health Affairs* 28(1):37–41.

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

29

Gold MR, Siegel JE, Weinstein MC, Russell LB (1996) *Cost-effectiveness in health and medicine* (New York, NY: Oxford University Press).

Gu S, Tang Z, Shi L, Sawhney M, Hu H, Dong H (2015) Cost-minimization analysis of metformin and acarbose in treatment of type 2 diabetes. *Value in Health Regional Issues* 6:84–88.

Guo R, Cheng L, Li J, Hahn PR, Liu H (2020) A survey of learning causality with data: Problems and methods. *ACM Computing Surveys* 53(4):1–37.

Hastie T, Tibshirani R, Friedman JH (2009) *The elements of statistical learning: Data mining, inference, and prediction* (New York NY: Springer), 2nd edition.

Helm JE, AhmadBeygi S, van Oyen MP (2011) Design and analysis of hospital admission control for operational effectiveness. *Production and Operations Management* 20(3):359–374.

Helm JE, Alaeddini A, Stauffer JM, Bretthauer KM, Skolarus TA (2016) Reducing hospital readmissions by integrating empirical prediction with resource optimization. *Production and Operations Management* 25(2):233–257.

Hoerl AE, Kennard RW (1970) Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12(1):55–67.

Ibrahim R, Kucukyazici B, Verter V, Gendreau M, Blostein M (2016) Designing personalized treatment: An application to anticoagulation therapy. *Production and Operations Management* 25(5):902–918.

Imbens GW, Rubin DB (2015) *Causal inference in statistics, social, and biomedical sciences* (New York, NY: Cambridge University Press).

Jakubik J, Feuerriegel S (2022) Data-driven allocation of development aid toward Sustainable Development Goals: Evidence from HIV/AIDS. *Production and Operations Management* 31(6):2739–2756.

Kahn HS, Cheng YJ, Thompson TJ, Imperatore G, Gregg EW (2009) Two risk-scoring systems for predicting incident diabetes mellitus in U.S. adults age 45 to 64 years. *Annals of Internal Medicine* 150(11):741–751.

Kaiser Permanente (2021) Annual report. URL https://about.kaiserpermanente.org/who-we-are/annual-reports/2021-annual-report.

Kamalzadeh H, Ahuja V, Hahsler M, Bowen ME (2021) An analytics-driven approach for optimal individualized diabetes screening. *Production and Operations Management* 30(9):3161–3191.

KC DS, Scholtes S, Terwiesch C (2020) Empirical research in healthcare operations: Past research, present understanding, and future opportunities. *Manufacturing & Service Operations Management* 22(1):73–83.

30

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

Keskinocak P, Savva N (2020) A review of the healthcare-management (modeling) literature published in Manufacturing & Service Operations Management. *Manufacturing & Service Operations Management* 22(1):59–72.

Khademi A, Saure DR, Schaefer AJ, Braithwaite RS, Roberts MS (2015) The price of nonabandonment: HIV in resource-limited settings. *Manufacturing & Service Operations Management* 17(4):554–570.

Knowler WC, Barrett-Connor E, Fowler SE, Hamman RF, Lachin JM, Walker EA, Nathan DM (2002) Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. *The New England Journal of Medicine* 346(6):393–403.

Koh HK, Sebelius KG (2010) Promoting prevention through the affordable care act. *The New England Journal of Medicine* 363(14):1296–1299.

LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521(7553):436–444.

Lee EK, Wei X, Baker-Witt F, Wright MD, Quarshie A (2018) Outcome-driven personalized treatment design for managing diabetes. *Interfaces* 48(5):422–435.

Lindstroem J, Tuomilehto J (2003) The diabetes risk score: A practical tool to predict type 2 diabetes risk. *Diabetes Care* 26(3):725–731.

Liu X, Hu M, Helm JE, Lavieri MS, Skolarus TA (2018) Missed opportunities in preventing hospital readmissions: Redesigning post-discharge checkup policies. *Production and Operations Management* 60(5):1019.

Long MT, Fox CS (2016) The framingham heart study: 67 years of discovery in metabolic disease. *Nature Reviews Endocrinology* 12(3):177–183.

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems* .

May JH, Spangler WE, Strum DP, Vargas LG (2011) The surgical scheduling problem: Current research and future opportunities. *Production and Operations Management* 20(3):392–405.

National Institute for Health and Care Excellence (2015) Medicines optimisation .

Niculescu-Mizil A, Caruana R (2005) Predicting good probabilities with supervised learning. *International Conference on Machine Learning.*

OECD (2019) Health at a glance 2019. URL `https://www.oecd-ilibrary.org/content/publication/4dd50c09-en`.

Paulweber B, Valensi P, Lindström J, Lalic NM, Greaves CJ, McKee M, Kissimova-Skarbek K, et al. (2010) A European evidence-based guideline for the prevention of type 2 diabetes. *Hormone and Metabolic Research* 42:3–36.

Rosella LC, Manuel DG, Burchill C, Stukel TA (2011) A population-based risk algorithm for the development of diabetes: Development and validation of the diabetes population risk tool DPoRT. *Journal of Epidemiology and Community Health* 65(7):613–620.

Rubin DB (2005) Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association* 100(469):322–331.

Schallmoser S, Zueger T, Kraus M, Saar-Tsechansky M, Stettler C, Feuerriegel S (2023) Machine learning for predicting micro-and macrovascular complications in individuals with prediabetes or diabetes: Retrospective cohort study. *Journal of Medical Internet Research* 25:e42181.

Senoner J, Netland T, Feuerriegel S (2022) Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science* 68(8):5557–6354.

Shi C, Blei D, Veitch V (2019) Adapting neural networks for the estimation of treatment effects. *Advances in Neural Information Processing Systems* .

Tachkov K, Mitov K, Koleva Y, Mitkova Z, Kamusheva M, Dimitrova M, Petkova V, et al. (2020) Life expectancy and survival analysis of patients with diabetes compared to the non diabetic population in Bulgaria. *PLOS ONE* 15(5):e0232815.

Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society* 58(1):267–288.

US Centers for Disease Control and Prevention (2020) National diabetes statistics report. URL `https://www.cdc.gov/diabetes/data/statistics/statistics-report.html`.

van der Laan MJ, Rubin D (2006) Targeted maximum likelihood learning. *The International Journal of Biostatistics* 2(1):Article 11.

Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.

Wilson PWF, Meigs JB, Sullivan L, Fox CS, Nathan DM, D'Agostino RB (2007) Prediction of incident diabetes mellitus in middle-aged adults: The Framingham offspring study. *Archives of Internal Medicine* 167(10):1068–1074.

32

Author: *Data-Driven Allocation of Preventive Care*

Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

World Health Organization (2006) Definition and diagnosis of diabetes mellitus and intermediate hyperglycaemia: Report of a WHO/IDF consultation.

World Health Organization (2016) Voluntary health insurance in europe: role and regulation .

World Health Organization (2018a) Global health estimates 2016: Disease burden by cause, age, sex, by country and by region, 2000-2016.

World Health Organization (2018b) Noncommunicable diseases. URL `http://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases`.

World Health Organization (2018c) The top 10 causes of death. URL `https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death`.

Zhang DJ, Gurvich I, Van Mieghem JA, Park E, Young RS, Williams MV (2016) Hospital readmissions reduction program: An economic and operational analysis. *Management Science* 62(11):3351–3371.

Zhou X, Siegel KR, Ng BP, Jawanda S, Proia KK, Zhang X, Albright AL, et al. (2020) Cost-effectiveness of diabetes prevention interventions targeting high-risk individuals and whole populations: A systematic review. *Diabetes Care* 43(7):1593–1616.

Zueger T, Schallmoser S, Kraus M, Saar-Tsechansky M, Feuerriegel S, Stettler C (2022) Machine learning for predicting the risk of transition from prediabetes to diabetes. *Diabetes Technology & Therapeutics* 24(11):842–847.

Online Supplement

Appendix A: Link of Prevented Diseases to Cost Minimization

Proof of Proposition 1. Without loss of generality, we study year l within our study period, and let $\{1, \ldots, k\} \subset \{1, \ldots, N\}$ denote the patients that were enrolled for preventive treatment according to our decision model, i.e., which describe a solution to the decision problem (Equation (6)). For simplicity, we neglect the index denoting the year. Then, the first part of this proof shows that the solution $\{1, \ldots, k\}$ minimizes costs among all subsets containing k patients. The second part proves that there is no subset with fewer patients $(\{1, \ldots, k-1\} \subset \{1, \ldots, N\})$ that yields lower costs.

1st Part. Because $\{1, \ldots, k\} \subset \{1, \ldots, N\}$ is a solution to the decision problem, we can assume that

$$\mathbb{E}\left[y \mid x_i, t_i = 0\right] - \mathbb{E}\left[y \mid x_i, t_i = 1\right] \geq \mathbb{E}\left[y \mid x_j, t_j = 0\right] - \mathbb{E}\left[y \mid x_j, t_j = 1\right],$$

$$\forall i \in \{1, \ldots, k\} \, \forall j \in \{k+1, \ldots, N\}. \tag{12}$$

We take a random patient $i^* \in \{1, \ldots, k\}$ and a random patient $j^* \in \{k+1, \ldots, N\}$ and exchange the treatment (i.e., we provide preventive treatment to patient j^* but not to patient i^*). We compute the expected costs C_{after} after this exchange and compare it to the expected costs C_{before} before the exchange. The computation yields the following:

$$C_{\text{after}} - C_{\text{before}} \tag{13}$$

$$= C_{\text{diab}} \left(\sum_{\substack{i=1 \\ i \neq i^*}}^{k} \mathbb{E}\left[y \mid x_i, t_i = 1\right] + \sum_{\substack{j=k+1 \\ j \neq j^*}}^{N} \mathbb{E}\left[y \mid x_j, t_j = 0\right] + \mathbb{E}\left[y \mid x_{j^*}, t_{j^*} = 1\right] + \mathbb{E}\left[y \mid x_{i^*}, t_{i^*} = 0\right] \right)$$

$$+ k\,C_{\text{prevent}} - C_{\text{diab}} \left(\sum_{i=1}^{k} \mathbb{E}\left[y \mid x_i, t_i = 1\right] + \sum_{j=k+1}^{N} \mathbb{E}\left[y \mid x_j, t_j = 0\right] \right) - k\,C_{\text{prevent}} \tag{14}$$

$$= C_{\text{diab}} \left(\sum_{\substack{i=1 \\ i \neq i^*}}^{k} \mathbb{E}\left[y \mid x_i, t_i = 1\right] + \sum_{\substack{j=k+1 \\ j \neq j^*}}^{N} \mathbb{E}\left[y \mid x_j, t_j = 0\right] + \mathbb{E}\left[y \mid x_{j^*}, t_{j^*} = 1\right] + \mathbb{E}\left[y \mid x_{i^*}, t_{i^*} = 0\right] \right.$$

$$\left. - \sum_{i=1}^{k} \mathbb{E}\left[y \mid x_i, t_i = 1\right] - \sum_{j=k+1}^{N} \mathbb{E}\left[y \mid x_j, t_j = 0\right] \right) \tag{15}$$

$$= C_{\text{diab}} \left(\mathbb{E}\left[y \mid x_{i^*}, t_{i^*} = 0\right] - \mathbb{E}\left[y \mid x_{i^*}, t_{i^*} = 1\right] + \mathbb{E}\left[y \mid x_{j^*}, t_{j^*} = 1\right] - \mathbb{E}\left[y \mid x_{j^*}, t_{j^*} = 0\right] \right) \tag{16}$$

$$\geq 0 \tag{17}$$

2

Author: *Data-Driven Allocation of Preventive Care*
Article accepted by *Manufacturing & Service Operations Management*; manuscript no. MSOM-21-251

2nd Part. We now show that no subset comprising fewer patients ($\{1,\ldots,k-1\} \subset \{1,\ldots,N\}$) yields lower costs. For this computation, we assume that:

$$C_{\text{diab}}\, \mathbb{E}\left[y \,|\, x_i, t_i = 1\right] + C_{\text{prevent}} \leq C_{\text{diab}}\, \mathbb{E}\left[y \,|\, x_i, t_i = 0\right] \qquad \forall i \in \{1,\ldots,k\}. \tag{18}$$

That an exchange of patients between the treatment group and the non-treatment group leads to increased costs directly follows from part 1. Thus, without loss of generality, we consider subgroups $\{1,\ldots,k-1\}$ of our original solution $\{1,\ldots,k\}$ to the decision problem. We compute costs C_{k-1} when providing preventive care to $k-1$ patients and compare C_{k-1} to the original costs C_k:

$$C_{k-1} - C_k \tag{19}$$

$$= C_{\text{diab}} \left(\sum_{i=1}^{k-1} \mathbb{E}\left[y \,|\, x_i, t_i = 1\right] + \sum_{j=k}^{N} \mathbb{E}\left[y \,|\, x_j, t_j = 0\right] \right) + (k-1)\, C_{\text{prevent}}$$

$$\quad - C_{\text{diab}} \left(\sum_{i=1}^{k} \mathbb{E}\left[y \,|\, x_i, t_i = 1\right] + \sum_{j=k+1}^{N} \mathbb{E}\left[y \,|\, x_j, t_j = 0\right] \right) - k\, C_{\text{prevent}} \tag{20}$$

$$= C_{\text{diab}} \left(\sum_{i=1}^{k-1} \mathbb{E}\left[y \,|\, x_i, t_i = 1\right] + \sum_{j=k}^{N} \mathbb{E}\left[y \,|\, x_j, t_j = 0\right] - \sum_{i=1}^{k} \mathbb{E}\left[y \,|\, x_i, t_i = 1\right] - \sum_{j=k+1}^{N} \mathbb{E}\left[y \,|\, x_j, t_j = 0\right] \right) - C_{\text{prevent}} \tag{21}$$

$$= C_{\text{diab}} \left(\mathbb{E}\left[y \,|\, x_k, t_k = 0\right] - \mathbb{E}\left[y \,|\, x_k, t_k = 1\right] \right) - C_{\text{prevent}} \tag{22}$$

$$= C_{\text{diab}} \left(\mathbb{E}\left[y \,|\, x_k, t_k = 0\right] \right) - C_{\text{diab}} \left(\mathbb{E}\left[y \,|\, x_k, t_k = 1\right] \right) - C_{\text{prevent}} \tag{23}$$

$$\geq 0 \tag{24}$$

$\square$

Appendix B: Proof for Optimal Allocation

Proof of Proposition 2.   Let n_{onsets} denote the number of expected onsets. Assume $h_0^*(x_{i,l})$ to be a perfect model for $\mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 0\right]$ and assume $\gamma_{i,l}$ to be the true treatment effect for $i \in \{1, \ldots, N\}$. Let $\{1, \ldots, k\}$ be the subset of patients who are prescribed preventive care, and let $\{k+1, \ldots, N\}$ be the subset of patients who are not prescribed preventive care in year l. Let this allocation follow our described approach to allocating patients; thus,

$$(1 - \gamma_{i,l})\, h_0^*(x_{i,l}) \geq (1 - \gamma_{j,l})\, h_0^*(x_{j,l}), \qquad \forall i \in \{1, \ldots, k\}\, \forall j \in \{k+1, \ldots, N\}, \tag{25}$$

which is equal to

$$h_0^*(x_{i,l}) + \gamma_{j,l}\, h_0^*(x_{j,l}) \geq h_0^*(x_{j,l}) + \gamma_{i,l}\, h_0^*(x_{i,l}), \qquad \forall i \in \{1, \ldots, k\}\, \forall j \in \{k+1, \ldots, N\}. \tag{26}$$

In the following computation, we take a random patient $i^* \in \{1, \ldots, k\}$ and a random patient $j^* \in \{k+1, \ldots, N\}$ and exchange the treatment—that is, we provide preventive treatment to patient j^* but not to patient i^*. From this exchange, we arrive at the following:

$$n_{\text{onsets}} = \sum_{\substack{i=1 \\ i \neq i^*}}^{k} \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 1\right] + \sum_{\substack{j=k+1 \\ j \neq j^*}}^{N} \mathbb{E}\left[y \,|\, x_{j,l}, t_{j,l} = 0\right] + \mathbb{E}\left[y \,|\, x_{j^*,l^*}, t_{j^*,l^*} = 1\right] + \mathbb{E}\left[y \,|\, x_{i^*,l^*}, t_{i^*,l^*} = 0\right] \tag{27}$$

$$= \sum_{\substack{i=1 \\ i \neq i^*}}^{k} (1 - \gamma_{i,l})\, h_0^*(x_{i,l}) + \sum_{\substack{j=k+1 \\ j \neq j^*}}^{N} h_0^*(x_{j,l}) + (1 - \gamma_{j^*,l^*})\, h_0^*(x_{j^*,l^*}) + h_0^*(x_{i^*,l^*}) \tag{28}$$

$$\geq \sum_{\substack{i=1 \\ i \neq i^*}}^{k} (1 - \gamma_{i,l})\, h_0^*(x_{i,l}) + \sum_{\substack{j=k+1 \\ j \neq j^*}}^{N} h_0^*(x_{j,l}) + (1 - \gamma_{i^*,l^*})\, h_0^*(x_{i^*,l^*}) + h_0^*(x_{j^*,l^*}) \tag{29}$$

$$= \sum_{i=1}^{k} (1 - \gamma_{i,l})\, h_0^*(x_{i,l}) + \sum_{j=k+1}^{N} h_0^*(x_{j,l}) \tag{30}$$

$$= \sum_{i=1}^{k} \mathbb{E}\left[y \,|\, x_{i,l}, t_{i,l} = 1\right] + \sum_{j=k+1}^{N} \mathbb{E}\left[y \,|\, x_{j,l}, t_{j,l} = 0\right] \tag{31}$$

Thus, exchanging the treatments of patients i^* and j^* leads to the same or a larger number of expected onsets. Consequently, our original allocation minimizes the number of expected onsets.   □

Appendix C: Robustness Checks

C.1.   Choice of Machine Learning Model

For the analysis in the main paper, the machine learning in stage 2 of the decision model was set to gradient boosted decision trees (Friedman 2001). Gradient boosted decision trees belong to the category of tree ensemble methods, which have been applied successfully in other operational applications (e. g., Glaeser et al. 2019, Senoner et al. 2022). We now provide empirical evidence supporting this choice. For this purpose, we compare the performance of our decision model under different machine learning models in stage 2 while the rest of the decision model remains the same. Our comparison includes linear (lasso and ridge regression) and nonlinear (random forest and deep neural network) models. All estimation details are provided in Supplement G.

   Results are shown in Table 7. While the lasso performs well among the linear models, nonlinear machine learning models consistently outperform linear models. This may be because nonlinear models are better able to exploit complex patterns in the data, particularly with regard to between-patient heterogeneity. However, nonlinear models can also be prone to overfitting, which may explain why the deep neural network underperformed compared to the gradient boosted decision trees. Overall, the decision model using gradient boosted decision trees achieves the best performance.

Table 7   Performance comparison among different machine learning models.

	k=1,000		k=5,000		k=10,000	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Naïve baseline	53.438	6.319	250.422	8.889	489.331	11.988
	(0.330)	(0.235)	(1.698)	(0.258)	(2.162)	(0.249)
Lasso	152.124	7.489	520.490	12.726	811.841	16.800
	(0.840)	(0.218)	(2.401)	(0.234)	(3.684)	(0.261)
Ridge regression	150.495	7.409	514.812	12.586	803.021	16.617
	(0.831)	(0.216)	(2.374)	(0.231)	(3.644)	(0.258)
Random forest	160.270	7.892	548.875	13.425	855.942	17.717
	(0.885)	(0.230)	(2.531)	(0.246)	(3.880)	(0.276)
Deep neural network	155.636	7.670	535.133	13.109	833.792	17.276
	(0.860)	(0.224)	(2.467)	(0.240)	(3.766)	(0.268)
Gradient boosted decision trees	162.913	8.050	567.710	13.979	882.004	18.337
	(0.902)	(0.239)	(2.612)	(0.253)	(3.934)	(0.284)

Stated: mean performance (standard deviation in parentheses)

Note. Performance metrics for allocation of preventive care when using the respective machine learning model in stage 2 of our decision model. Budget constraints allow for the treatment of k patients per year. Cost savings are reported in USD millions.

C.2.   Comparison with Alternative Risk Scores

The Framingham diabetes risk score represents the quasi-standard in clinical practice for assessing diabetes risk (Long and Fox 2016). Nevertheless, Table 8 shows the comparison between our decision model and other diabetes

risk scores from clinical practice—namely, Lindstroem and Tuomilehto (2003), Wilson et al. (2007), Kahn et al. (2009), and Rosella et al. (2011). Consistent with our previous findings, our proposed decision model outperforms all alternative risk scores.

Table 8   Performance comparison between our decision model and other risk scores.

	k=1,000		k=5,000		k=10,000	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Naïve baseline	53.438	6.319	250.422	8.889	489.331	11.988
	(0.330)	(0.235)	(1.698)	(0.258)	(2.162)	(0.249)
Lindstroem and Tuomilehto (2003)	128.810	6.323	434.431	10.562	679.740	14.015
	(0.710)	(0.181)	(2.008)	(0.196)	(3.126)	(0.219)
Wilson et al. (2007)	121.424	5.979	415.914	10.173	648.571	13.425
	(0.671)	(0.174)	(1.918)	(0.187)	(2.940)	(0.209)
Kahn et al. (2009)	119.542	5.879	406.948	9.930	635.441	13.133
	(0.660)	(0.170)	(1.878)	(0.183)	(2.897)	(0.205)
Rosella et al. (2011)	120.917	5.940	409.335	9.966	639.950	13.207
	(0.667)	(0.171)	(1.891)	(0.184)	(2.932)	(0.206)
Our decision model	162.913	8.050	567.710	13.979	882.004	18.337
	(0.902)	(0.239)	(2.612)	(0.253)	(3.934)	(0.284)

Stated: mean performance (standard deviation in parentheses)

Note. Performance metrics for allocation of preventive care with the respective decision model. Budget constraints allow for the treatment of k patients per year. Cost savings over no preventive care allocation are reported in USD millions.

C.3. Robustness to Estimation Errors in the Treatment Effect

We now analyze the robustness of our decision model to potential errors in the estimation of the treatment effect. For this, we run the same analysis as before, but introduce additional Gaussian noise to the estimated treatment effect $\gamma_{i,l}$, i.e.,

$$\tilde{\gamma}_{i,l} = \gamma_{i,l} + \epsilon \qquad \text{with} \qquad \epsilon \sim \mathcal{N}(0, \sigma^2), \tag{32}$$

where σ denotes the level of noise. The latter thus represents the estimation error. We further know that the effectiveness of preventive care such as *metformin* is bounded and, therefore, ensure the noisy treatment effect $\tilde{\gamma}_{i,l}$ to lie in the interval $[0, 1]$ through clipping.

Table 9 lists the results for this analysis. Here, we compare different noise levels $\sigma \in \{0, 0.1, 0.5\}$. Of note, a noise level of $\sigma = 0.5$ introduces a large estimation error, especially when considering that the effectiveness of preventive treatments should lie in the interval $[0, 1]$ for real-world clinical settings. As expected, we find a slightly more dominant role of estimation errors for small k, which, in light of the small set of patients allocated to preventive care, results in slightly reduced cost savings. This can be expected due to the small sample size. However, and more importantly, we find that the impact of estimation errors is overall fairly small: we still achieve a large number of prevented disease onsets as well as large cost savings over current practice. For example, even

for noise of $\sigma = 0.1$, our decision model still outperforms current practice for $k = 10,000$ and is only slightly outperformed by current practice for noise of $\sigma = 0.5$.

Table 9   Robustness of our decision model to estimation errors in the treatment effect.

	$k=1,000$		$k=5,000$		$k=10,000$	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Clinical baseline	134.798	7.340	446.948	11.996	701.942	15.350
	(1.025)	(0.245)	(3.373)	(0.258)	(3.837)	(0.265)
$\sigma = 0$	162.913	8.050	567.710	13.979	882.004	18.337
	(0.902)	(0.239)	(2.612)	(0.253)	(3.934)	(0.284)
$\sigma = 0.1$	149.481	7.331	501.655	12.172	785.779	16.180
	(0.823)	(0.209)	(2.320)	(0.227)	(3.630)	(0.253)
$\sigma = 0.5$	131.199	6.462	449.977	11.012	701.492	14.526
	(0.725)	(0.189)	(2.075)	(0.202)	(3.176)	(0.226)

Stated: mean performance (standard deviation in parentheses)

Note. Budget constraints allow for the treatment of k patients per year. Cost savings over no preventive care allocation are reported in USD millions.

C.4. Robustness to Number of Training Samples

In machine learning, the size of the training set is known to affect the prediction performance, which, in turn, affects the operational performance. We thus study the convergence of our decision model against the optimal decision for an increasing number of observations. We use a static setting with observations $x \in \mathbb{R}^3$, for which the three components represent different patient variables (age, height, body mass index). We simulate these as follows:

$$x \sim \mathcal{N}(\mu, \Sigma), \tag{33}$$

$$\mu = (50, 170, 27)^T, \tag{34}$$

$$\Sigma = \begin{bmatrix} 20 & 0 & 5 \\ 0 & 50 & 5 \\ 5 & 5 & 5 \end{bmatrix} \tag{35}$$

The true but unobservable risk of an onset is given by

$$y = \sigma \left(\frac{0.5\, x_1 + 0.1\, x_2 + 0.2\, x_3}{100} \right), \tag{36}$$

with σ denoting the sigmoid function. We assume that we observe the actual outcomes only for N_{train} patients. That is, for each patient i, we assume that we observe the realization of Y given by

$$y_i = \begin{cases} 1, & \text{if } y + \epsilon \geq 0.7, \\ 0, & \text{otherwise,} \end{cases} \quad \text{with} \quad \epsilon \sim \mathcal{N}(0, 1). \tag{37}$$

In the simulation, we further set the true treatment effect to $\gamma = 0.31$.

Our population comprises $N = 100{,}000$ patients. The budget constraint allows us to enroll $k = 10{,}000$ patients into preventive care. Figure 4 shows the number of prevented diseases for varying training sample sizes N_{train}. It confirms that our decision model converges to the optimal allocation. Moreover, we already yield a close-to-optimum performance for a relatively small number of training samples (i. e., fewer than 3,000).

Figure 4   **Performance of our decision model for varying number of training samples.**
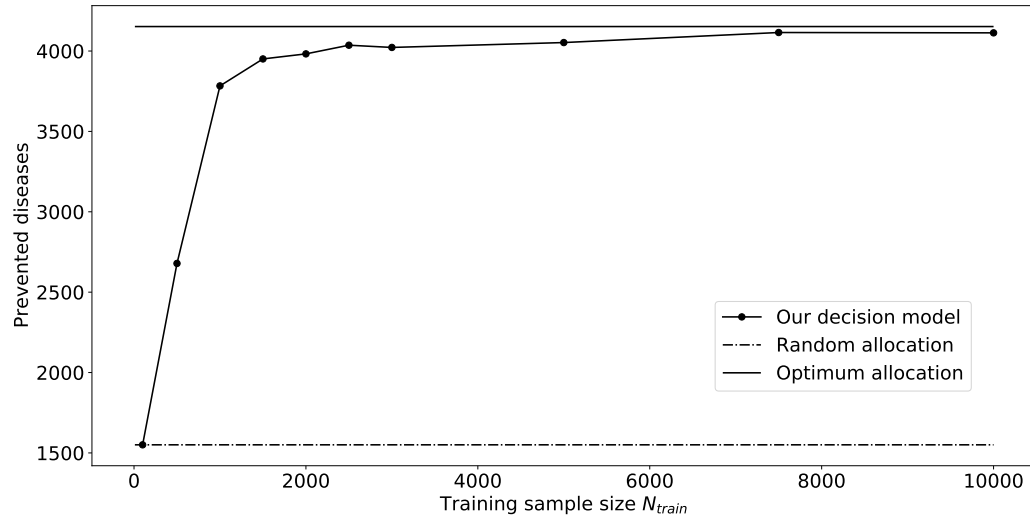
Appendix D: Causal Sensitivity Analysis

In the following, we analyze how potential unobserved confounders would affect the estimated treatment effects. Unobserved confounders refer to unmeasured variables that affect both treatment assignment and the health outcome (Cinelli and Hazlett 2020). A prominent example is the socioeconomic status of patients, as wealthier patients are often more likely to enroll in preventive care and, furthermore, can afford a healthier lifestyle that is responsible for a lower diabetes risk, which thus may underestimate the treatment effect of *metformin*. Since these confounders are unobserved, we cannot control for them when estimating the effect of a treatment. As a remedy, we follow the so-called causal sensitivity analysis from Cinelli and Hazlett (2020) and analyze the effect that an unobserved confounder would have on the treatment effect estimation, if it existed. More precisely, we select an observed variable which affects both treatment assignment and the health outcome and now assume that an unobserved confounder exists with the same strength as the observed one. Then, we check that the treatment effect cannot be explained away and thus remains robust.

The age of a patient is generally associated with the likelihood of being prescribed with metformin (e.g., Rosella et al. 2011) and has also been shown to be strongly associated with the risk of diabetes onset (e.g., Knowler et al. 2002). Thus, we use the strength of the variable "age" in the following as our reference to make comparisons of whether the treatment effect can be explained away.

We follow Cinelli and Hazlett (2020) and use a linear regression. Informed by clinical research (e.g., the Framingham risk score), we use a linear regression with the following patient variables, i.e.,

$$y = treatment + age + sex + weight + body\_mass\_index + systolic\_bp + diastolic\_bp + height. \tag{38}$$

We then perform causal sensitivity analysis for all patients and first test for the sensitivity of the main effect. Afterward, to account for heterogeneity in the treatment effect, we repeat the causal sensitivity analysis for different patient subgroups. For this, we segment the patients by their estimated treatment effects into four groups – A, B, C, D – using a decision tree regressor, and then perform the causal sensitivity analysis for each subgroup.

Table 10 shows the results of our causal sensitivity analysis. The results confirm that the estimated treatment effect is robust to unobserved confounders. The results are also consistent across all subgroups. Note that the patient's age is considered to be one of the strongest predictors for diabetes (American Diabetes Association 2022), and, thus, it is very unlikely to have an unobserved confounder of a similar strength. In other words,

the treatment effect cannot be explained away by unobserved confounding. In sum, we do not see evidence that unobserved confounders may undermine our estimates of treatment effectiveness, which is in line with many existing works using randomized control trials that have already confirmed the effectiveness of *metformin* (e.g., Knowler et al. 2002).

Table 10   **Results of robustness check for unobserved confounders.**

Strength	All patients	Subgroups			
		A	B	C	D
$0.2 \times age$	0.066	0.066	0.142	0.048	0.058
	[0.052, 0.088]	[0.045, 0.096]	[0.103, 0.188]	[0.026, 0.074]	[0.027, 0.097]
$0.5 \times age$	0.053	0.052	0.127	0.037	0.046
	[0.039, 0.074]	[0.029, 0.078]	[0.084, 0.177]	[0.015, 0.060]	[0.019, 0.075]
$0.8 \times age$	0.043	0.040	0.117	0.029	0.039
	[0.031, 0.058]	[0.013, 0.072]	[0.075, 0.132]	[0.009, 0.052]	[0.014, 0.071]

Stated: mean estimation (95% confidence intervals in parentheses)

Note. Estimated treatment effects when assuming the existence of an unobserved confounder with the strength $0.2 \times age$, $0.5 \times age$, and $1.0 \times age$.
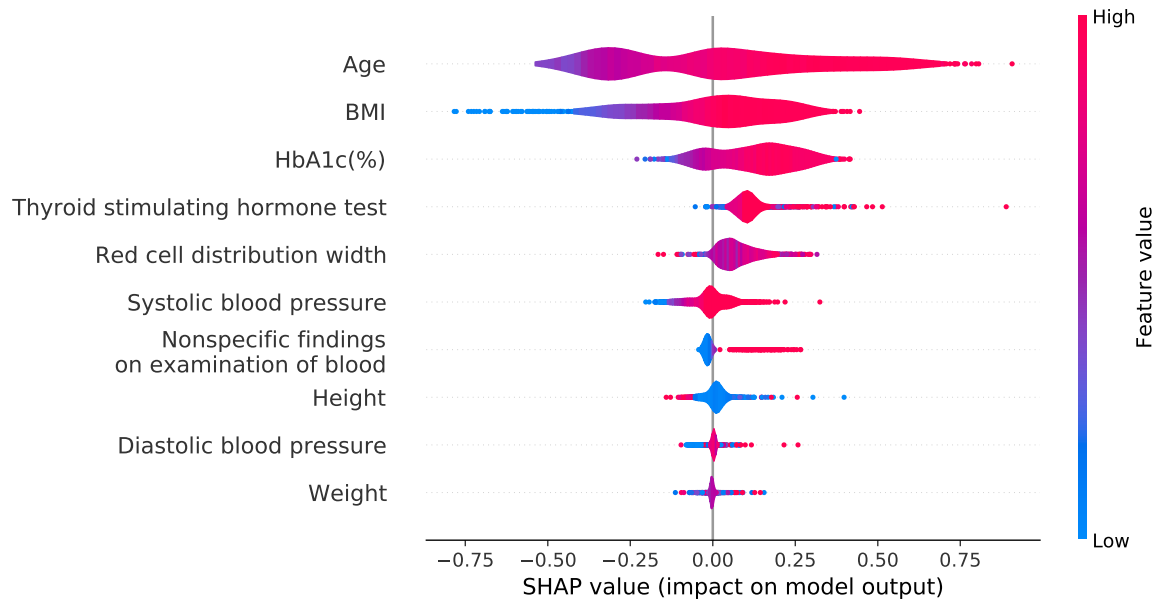
Appendix E: Machine Learning Explainability

To determine the most influential predictors, we calculated SHAP values. SHAP values are a unified method for measuring how much a predictor contributes to the overall model output and, therefore, rank the importance of the features. Moreover, they indicate whether higher (lower) values of a predictor are associated with an increased risk of diabetes onset.

Figure 5 shows the ten most important predictors for the gradient boosted decision tree when predicting diabetes onset. The most important predictors are age, body mass index (BMI), and the current glycated hemoglobin level (HbA1c). Age and BMI are well known risk factors for diabetes, and both were accordingly among the most important predictors for our gradient boosted decision tree (Abbasi et al. 2016). Also among the ten most important predictors are the red cell distribution width and thyroid stimulating hormone tests, which have both been found in the medical literature to be predictors for diabetes (e.g., Chaker et al. 2016). The height and weight of the patient are also among the ten most important predictors, yet with less importance. This is likely due to the fact that the predictive power of these variables as individual predictors is limited, and only their combination (as in the BMI) leads to an important predictor.

We discussed our results with medical experts from diabetes care, who confirmed to us that these predictors are well-established risk factor for diabetes.

Figure 5   **SHAP values for the ten most important predictors for the prediction model. The dots represent observation**
for which the x-axis denotes the SHAP value. Observations with a positive (negative) SHAP value denote
a higher (lower) risk for a diabetes onset. The color of each dot indicates the corresponding value of the
predictor for a given observation.

Appendix F: Calibration of Machine Learning Model

We use two techniques to improve the calibration of our machine learning model. Here, the aim is that the distribution of the predicted probability is similar to the distribution of the observed probability in the training data. First, we use the synthetic minority over-sampling technique (SMOTE) (Chawla et al. 2002) to oversample observations from patients who develop diabetes. Second, we use Platt scaling to calibrate the predicted class probabilities (Niculescu-Mizil and Caruana 2005). Figure 6 shows the calibration plot for our gradient boosted decision trees after resampling. The calibration plot shows a good calibration.

Figure 6   **Calibration of our machine learning model. The gray area shows the standard deviation due to resampling.**
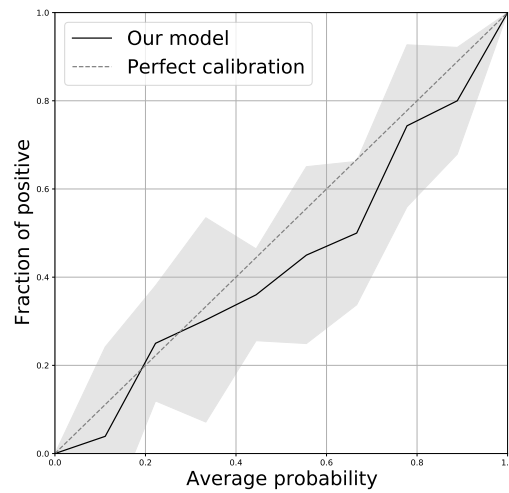
Appendix G: Hyperparameter Tuning

We determined the hyperparameters using Bayesian optimization with 100 iterations of optimization. Table 11 lists the tuning ranges for each parameter. In the Bayesian optimization, the initial parameters are sampled from a uniform distribution. We implemented and evaluated all decision models in Python. The lasso, ridge regression, and random forest machine learning models are implemented using scikit-learn. The deep neural network is implemented using Pytorch. We train the deep neural network using the Adam optimizer. We use ReLU activation in the hidden layers and sigmoid in the output layer. The objective of the optimizer is to minimize the binary cross-entropy. The gradient boosted decision trees are implemented using LightGBM. For the causal forest, we keep all parameters at their default values. That is, we set the number of estimators to 10, the number of features to consider when looking for the best split to 10, the maximum depth of the trees to 5, the minimum number of samples required to be split at a leaf node to 100, and the minimum number of samples required of the experiment group to be split at a leaf node to 10.

Table 11   **Hyperparameter tuning.**

Model	Tuning parameters	Tuning range
Lasso	Regularization strength α	$(0.001, 10000)$
Ridge regression	Regularization strength	$(0.001, 10000)$
Random forest	Number of trees	$(20, 200)$
	Min samples for split	$(2, 150)$
	Number of features to consider for split	{sqrt, log2, all}
	Criterion for split	{gini, entropy}
	Class weight	{none, balanced}
Deep neural network	Number of neurons	$(5, 128)$
	Number of hidden layers	$(1, 5)$
	Regularization strength	$(0.0, 5.0)$
	Learning rate	$(0.0001, 0.1)$
	Batch size	$(10, 1024)$
Gradient boosted decision trees	Number of trees	$(20, 200)$
	Number of leaves	$(20, 150)$
	Learning rate	$(0.01, 0.5)$
	Number of samples for constructing bins	$(20000, 300000)$
	Minimum number of data needed in a child	$(20, 500)$
	L1 regularization term on weights	$(0, 1)$
	L2 regularization term on weights	$(0, 1)$
	Class weight	{none, balanced}

Appendix H: Applicability for Settings with Known Treatment Effects

Here, we aim to demonstrate the applicability of our model in settings where the treatment effect of preventive care is known. In such a case, there is no need for healthcare organizations to perform counterfactual inference for treatment effect estimation, as the treatment effect can the be directly entered in the decision model.

To show the applicability of our decision model, we build upon a different setting. Recall that we examined *metformin* in the main paper. The reason was that it represents the quasi-standard for preventive care aimed at individuals at risk of developing diabetes in current medical practice (Long and Fox 2016). Several studies also suggest a risk reduction if patients are enrolled in personalized coaching toward a healthier lifestyle; however, such personalized coaching is currently not covered by health insurers in many countries (e.g., Germany, Switzerland, and Medicare in the United States). Nevertheless, it may be interesting to study the cost-effectiveness of lifestyle coaching for diabetes prevention.

In the case of personalized coaching, the goal of preventive treatment is to achieve lifestyle changes of the patients. Lifestyle changes have not been prescribed to the customers of our partnering health insurer. Thus, the treatment effect cannot be estimated through stage 1 of our decision model. Instead, we rely on prior literature that has estimated the treatment effect of personalized coaching through randomized controlled trials (Knowler et al. 2002).

We assume a treatment effect due to personalized coaching of $\gamma_{i,l^*} = 0.58$ (Knowler et al. 2002) that is exponentially decreasing over time, i.e.,

$$\gamma_{i,l} = \gamma_{i,l^*+1} \exp(l^* - l), \quad \text{for } l \geq l^*. \tag{39}$$

We set the annual cost of personalized coaching to $C_{\text{prevent}} = 1,600$ (Azelton et al. 2021).

Table 12 shows the empirical results. Here, we again compare the data-driven allocations from our decision model against that from current practice. Our findings are line with the results from the main paper: our decision model outperforms current practice across all metrics.

To sum up, there are different benefits of whether the treatment effect is (a) known or (b) estimated through stage 1 of our decision model. A benefit of (a) is that healthcare organizations can leverage known treatment effects from randomized controlled trials, which represent the gold standard for measuring treatment effects. In particular, there is no bias due to unobserved confounders. A benefit of (b) is that randomized controlled trials are not always available or otherwise costly, because of which our estimation may be preferred in practice. Moreover, using counterfactual inference in our decision model offers a mathematical approach to directly

learn heterogeneous treatment effects and thus account for the differential effectiveness of preventive care across

patients.

Table 12   **Comparison of current practice vs. our data-driven decision model for the case of personalized coaching.**

	k=1,000		k=5,000		k=10,000	
	Prevented diseases	Cost savings	Prevented diseases	Cost savings	Prevented diseases	Cost savings
Naïve baseline	728.378	16.457	1423.607	26.651	2736.326	45.807
	(5.606)	(0.141)	(10.761)	(0.249)	(16.553)	(0.323)
Current practice	1318.303	25.786	2083.693	37.088	3184.516	51.263
	(9.893)	(0.223)	(12.719)	(0.287)	(14.837)	(0.310)
Our decision model	1555.792	29.544	2143.867	38.036	3349.127	53.767
	(10.960)	(0.223)	(14.553)	(0.295)	(15.103)	(0.276)

Stated: mean performance (standard deviation in parentheses)

Note. Performance metrics for allocating preventive care given a varying budget for enrolling k patients per year into preventive treatments. Cost savings over no preventive care allocation are reported in USD millions.